KM.GOV

Semantic Interoperability Community of Practice (SICoP)

Semantic Wave 2006:
Executive Guide to the Business Value of Semantic Technologies

White Paper Series Module 2

Updated on 01/9/06

Version 1.0

SICoP White Paper Series Module 2

Semantic Wave 2006: Executive Guide to Billion Dollar Markets

Executive Editors and Co-Chairs

Dr. Brand Niemann, *U.S. EPA, Office of the CIO (SICoP Co-Chair)*

Harriet J. Riofrio, *Senior Staff Officer for Knowledge Management, Office of Assistant Secretary of Defense for Networks and Information Management, Deputy Chief Information Officer, Information Management (OASD NII DCIOIM), U.S. Department of Defense (KM.Gov Co-Chair)*

Earl Carnes, *Nuclear Industry Liaison, Environment, Safety & Health, Office of Regulatory Liaison, U.S. Department of Energy (KM.Gov Co-Chair)*

Principal Investigator

Mills Davis,
Project10X
202-667-6400
mdavis@project10x.com

Reviewers

Harriet RioFrio
OASD/NII DCIOIM
703-602-0516
Harriet.Riofrio@osd.mil

George Strawn
NSF, CIO
703-292-8100
gstrawn@nsf.gov

Michael Novak
IRS, Office of Research
202-874-0335
Michael.j.novak@irs.gov

Michael Todd
OASD/NII, CIO Staff
703-602-0516
Michael.Todd@osd.mil

Semantic Wave 2006

Part-1: Executive Guide to Billion Dollar Markets

A Project10X Special Report

January 2006

Project10X
2853 Ontario Road NW #501
Washington DC 20009 USA

Mills Davis
Managing Director
mdavis@project10x.com
202-667-6400

www.project10x.com

Contents

Figures

1. About this guide

What is the purpose and scope of this guide?

Semantic Wave 2006—Part-1 is a quick guide to the business value of semantic wave technologies and markets. It provides information and perspectives that are intended to help: (a) business line executives seeking strategic advantage in global market spaces by exploiting next generation technology; (b) technologists and integrators responsible for R&D of advanced capabilities; and (c) venture capitalists and entrepreneurs seeking to pioneer high-growth business opportunities.

The guide is organized as follows:

- **Semantic wave** — Introduces "distributed intelligence" as a long wave of investment involving fundamental shifts in paradigm, technology, and economics. Depicts the cumulative economics of the semantic wave.
- **Semantic technologies** — Defines semantic technologies. Highlights trends and drivers that are shaping this semantic wave, and the direction it is heading. Explains what aspects of information and communication technologies (ICT) they impact.
- **Business value of semantic technologies** — Discusses the business value of semantic technologies. Describes capabilities of semantic technologies and how these impact development, infrastructure, information and knowledge, information-intensive applications, knowledge-intensive applications, system behaviors, and intellectual property. Explains how semantic technologies improve lifecycle return on investment (ROI). Summarizes performance of semantic technologies in terms of efficiency, effectiveness, and strategic edge.
- **Signs of life in the marketplace** — Overviews early adopter case examples illustrate where, how, and in what ways semantic technologies are being applied, and to what effect. Identifies companies that are working on semantic technology R&D, products, and service delivery.
- **Semantic wave market view** — Characterizes semantic wave markets, their structure, size and growth, and how they will evolve to 2015. Sets forth the basis for this outlook.

2. Semantic Wave

Figure-01 **Long waves of innovation**

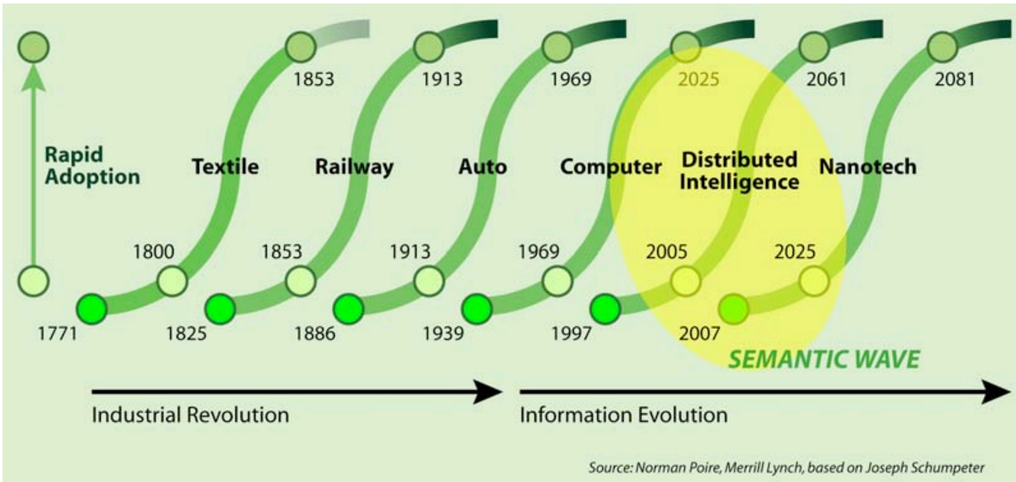

What is the semantic wave?

The semantic wave is a long wave of investment involving fundamental shifts in paradigm, technology, and economics. It will provide infrastructure and reasoning engines to fuel exponential economic expansion.

Looking back over the past two centuries, major conceptual advances that power economic growth seem to occur about twice a century. Today we are at the intersection of three major innovation advances: one nearing its end; one that will continue another 20-30 years; and one that is just starting. These long waves of innovation spur enormous investments and radically alter the economics of affected industries. As with the computer wave, the current one "distributed intelligence" is affecting virtually all industries.

Joseph Schumpeter, an Austrian-born economist, noted long waves of industrial activity in the 1940s. More recently, Merrill Lynch analyst Norman Poire sketched out a diagram (see Figure-01) that illustrates Schumpeter's concept. We've added a "you are here" overlay to Poire's diagram to indicate the current intersection of waves of innovation that comprise the "semantic wave". One upgrade to Schumpeter from thinkers such as Ray Kurzweill is that the frequency and amplitude of these long waves is accelerating. In fact, the evolutionary curve is not linear, but exponential.

When traced back to the Industrial Revolution in 18th-century England, Schumpeter noticed that waves of innovation ebbed and flowed every 50–60 years. Each fresh wave had brought with it a "new economy" that led to investment and excess, followed by a shakeout—but, ultimately, as The Economist concluded, left the world a richer and better place ("A Crunch of Gears," Economist, Sept. 29, 2001).

The chart shows six long waves. Inventions in cotton- spinning, iron-making, and steam power propelled the first boom. It lasted from the 1780s to the 1840s. The second wave arrived with innovations in steel-making and railways, lasting for half a century before running out of steam around 1900. Electrification and the internal-combustion engine powered the third 50- year wave. The fourth industrial wave was launched in the early 1950s on the back of petrochemicals, electronics, computing and aerospace. The fifth wave, distributed intelligence, started in the 1970s with the precursors of the Internet. It continued with the adoption of client-server corporate networking, and rapidly accelerated following the introduction of the World Wide Web. In the wake of the dot-com shakeout, this wave is shifting into a new growth gear. That's right: Far from being over, the current wave has probably another 35 years to go. Meanwhile, a sixth wave is forming that will be powered by nanotechnology, bioscience and clean energies.

New surges of economic activity tend to play out in four distinct phases. The first phase is a period of rapid innovation as practical applications of seminal inventions emerge. The next phase brings rapid growth as successful participants—whether in cotton, railways, motorcars, electrical goods or petrochemicals— enjoy fat margins, set standards, kill off weaker rivals and establish themselves as leaders of the pack. (In the information and communications technology [ICT] space, we might think of Cisco, Intel and Microsoft as leaders today; but will they continue their dominance during the next wave?) In the third phase, the market matures and the dominant firms hunker down for slower growth, which is happening now with the PC. The final phase is a short and sharp decline that occurs when the next set of technologies start jostling for the attention of investors.

What are the economics of the semantic wave?

The semantic wave will accelerate from vision and R&D to early adoption and mainstream markets valued in tens of billions of dollars by the end of this decade. The initiating force from the outset of this cycle has been public funding of fundamental R&D. The driving forces for mass-market adoption, however, are new capabilities combined with gains in efficiency, effectiveness, and other measures of performance.

Figure-02[1] depicts the cumulative flow of investment into and returns from semantic technology from 1995 to 2010. The three lines of investment are shown: (1) government R&D funding, (2) Venture capital investment and ICT company product/service commercialization, and (3) early adopter R&D. The two major lines of market return are: (a) returns (sales) for the technology vendors who've been commercializing the semantic technology; and (b) returns for the adopters of semantic technologies, which are measured in terms of gains in efficiency, effectiveness, and edge. Positive is up. Negative is down. Time moves left to right. So long as investment outlays exceed returns the line trends down. When returns start exceeding investments (and total cost of ownership) then the line moves upward.

Figure-02 Economics of the semantic wave

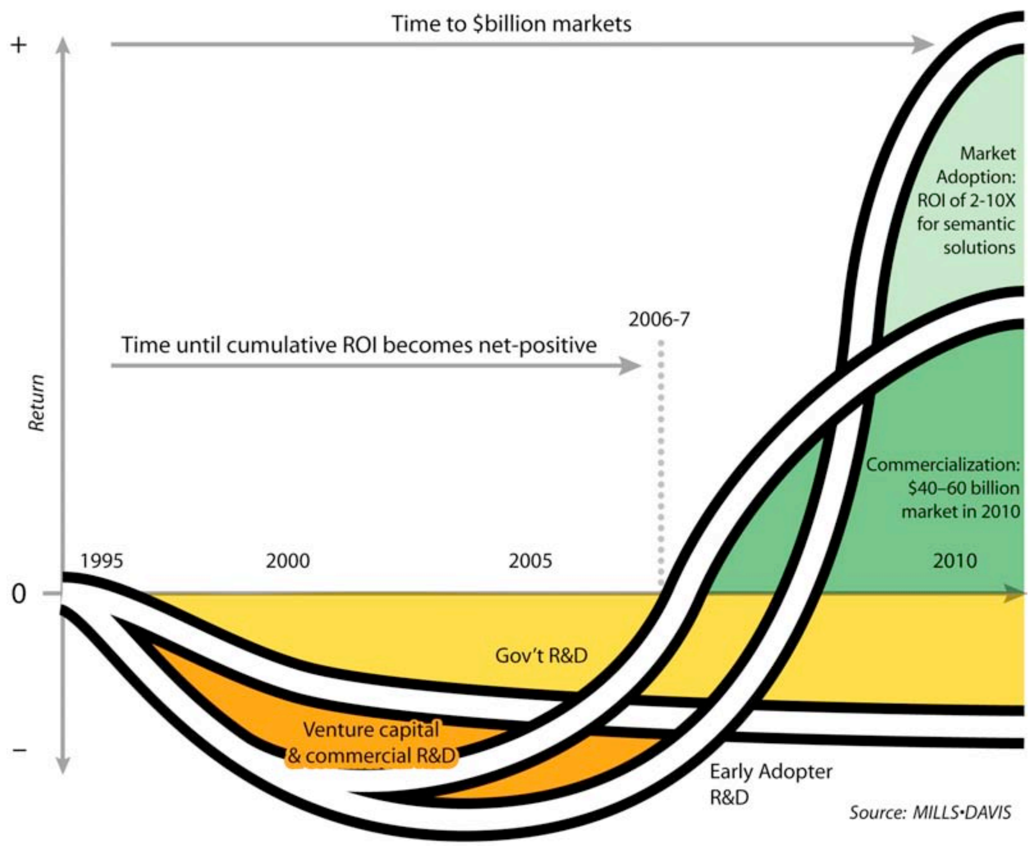

[1] Notes regarding market sizing, sources, and estimating methods will be provided in *Semantic Wave 2006: Part-2.*

3. Semantic Technologies

Why are semantic technologies needed now?

We face major challenges and opportunities that cannot be addressed successfully with contemporary information and communications technologies.

Challenge of net-centric infrastructure

Current ICT markets are dominated by relational database (RDB) and document-centric information technologies, procedural algorithmic programming paradigms, and stack architecture. The installed base and global market is huge — around $1.2T for hardware, software and services. Automation of transaction systems, proliferations of PCs, and global access through the worldwide web have been crowning achievements.

A key driver of global economic expansion in the coming decade is the build-out of broadband telecommunications and the deployment of intelligent services across this infrastructure.

The technology power curve driving this transformation has been widely recognized. Speed or capacity is doubling (or, more or less equivalently, price is halving) for network, storage and computing components every 9, 12 and 18 months, respectively.

Current hardware and software technologies have reached, or are rapidly approaching the limits of what they can do to cope with massive increases in scale, complexity, and lifecycle cost:

- *Scale* — explosion of infrastructure, information sources, communities of interest, and knowledge that characterizes todays net-centric global marketplace.
- *Complexity* — IT approaches inability to handle the integration and interoperability of systems and information.
- *Lifecycle cost* — Excessive total cost of ownership (TCO), where more than 70-percent of IT budgets must go to maintain steady-state system silos.

Achieving the gains in performance, agility, and lifecycle economics envisioned for net-centric environments demands architecture and technology solutions designed for the era of distributed intelligence, not for the desktop or the client-server world.

This vision of net-centric computing is often expressed as a "grand challenge" for the industry to develop systems and processes that are self-declaring, self-integrating, self-optimizing, self-protecting and self-healing and that can scale from point-to-point semantic web services to pervasive service grids. Key infrastructure challenges include security, pervasive services, stack complexity, parallelism, autonomic systems, legacy conversion, and application authoring. High-performance solutions are needed that exploit communications bandwidth and dynamic, massively distributed resources while minimizing the labor required to develop, integrate, maintain, and evolve such large scale computing and information environments.

Large un-met needs exist at the "hinges of the business" — integration and interoperability needs that go across systems, across businesses, and across communities of interest.

Conventional technologies have attempted to manage this complexity through layered abstraction. This method evolved from the early 1940's when simple theories, hardware, and primitive computing methods were designed as work-arounds for very expensive memory and storage resources. It made economic sense to continually re-compute rather then store sizable volumes of data.

Today, the requirements and economics of that early model have completely reversed. Theory and data requirements are now massive and complex and memory and storage costs are low. Yet, the conventions of repeat computation and structured data schemas have continued to the point where technology is bumping-up against physical and complexity limits that were unthinkable in the 40's.

Successful infrastructure solutions must resolve complexity and deliver business value at far less cost and risk than with prevailing approaches.

Challenge of information-intensive knowledge work

Knowledge work automation and knowledge worker augmentation are great un-met needs.

Businesses have already invested in enterprise applications that automate transaction processing. These "systems of record" enabled the business to scale. The problem is that maintenance of silo applications now consumes more than 70 percent of the IT budget. Current approaches do not give the leverage needed to reduce total cost of ownership (TCO) and free resources for innovation. Status quo is unsustainable.

Competitive differentiation in the industries is no-longer about automating transactions and record keeping, it is about enabling knowledgeable interactions with customers and suppliers, as well as across functions within the enterprise. It is about managing the processes and linkages between systems, handling exceptions to transactions, compliance, risk, fraud, and emergencies. In short, strategic advantage comes from integrating information and applications to automate knowledge worker functions. Stovepipe applications are too rigid, brittle, difficult to integrate, and prohibitively expensive to replace.

A new approach is needed to information-intensive knowledge work — one that automates capture of events, can connect the dots between people, places, and events using information from many different sources in different formats (structured & unstructured), followed by human monitoring and analysis of situations, workflows, and in-context collaboration and communication.

Successful solutions must be able to link applications, data sources, and services in easily used composite views, providing real-time interaction, analyisis, and decision-support.

Challenge of knowledge-computing

The challenge of knowledge work automation and knowledge worker augmentation involves more than making systems and information interoperable. It's about computing with knowledge — all the theory and information necessary to do the task.

Today's information systems focus on bringing information to the job, that is, situation awareness. But, the knowledge required to do a job is something an employee has to bring with him or her (via previous education and experience) or learn (on the job or by formal training). This education is expensive to acquire and provision. When people leave, the knowledge is rapidly lost to the organization.

To automate work on the farm, in the factory or in the office, the knowledge required to accomplish the task is laboriously hard-coded into mechanical parts, circuitry and software algorithms. Improvements in capability require repeated investments in next-generation solutions.

A new approach to knowledge-intensive work is needed that delivers not only the information, but all of the theory and modes of reasoning needed to perform a job or task.

Successful solutions will need to compute knowledge (both theory and information), plus handle very large-scale knowledgebases, complex forms of situation assessment, sophisticated value-based modes of reasoning, and autonomic and autonomous system behaviors. These challenges exceed the capabilities and performance capacity of current open standards approaches to knowledge representation and system architecture.

What are semantic technologies?

Semantic technologies are a new paradigm — an approach that deals with the challenges of net-centric infrastructure, knowledge work automation, and building systems that know what they're doing.

Semantic technologies are functional capabilities that enable both people and computers to create, discover, represent, organize, process, manage, reason with, present, share, and utilize meanings and knowledge to accomplish business, personal, and societal purposes.

Semantic technologies are tools that represent meanings, associations, theories, and know-how about the uses of things separately from data and program code. This knowledge representation is called an ontology — a run-time semantic model of information, defined using constructs for:

- Concepts – classes, things
- Relationships – properties (object and data)
- Rules – axioms and constraints
- Instances of concepts – individuals (data, facts)

Semantic models are like and unlike other IT models:

- Like databases, ontologies are used by applications at run time (queried and reasoned over). Unlike conventional databases, relationships are first-class constructs.
- Like object models, ontologies describe classes and attributes (properties). Unlike object models, ontologies are set-based and dynamic.
- Like business rules, semantic models encode event-based behaviors. Unlike business rules, ontologies organize rules using axioms.
- Like XML schemas, they are native to the web (and are in fact serialized in XML). Unlike XML schemas, ontologies are graphs not trees, and used for reasoning.

Figure-03 R&D themes in the semantic wave

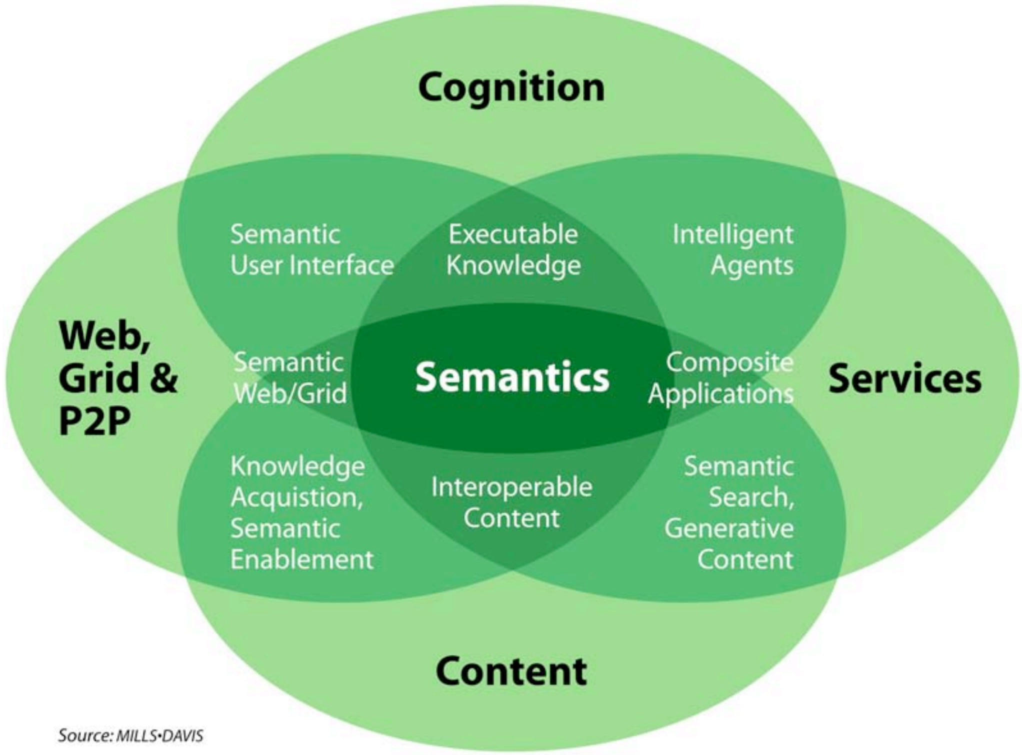

What is the scope of semantic technology R&D?

Semantic technologies have emerged as a central theme across a broad array of ICT research and development initiatives.

Figure-03 visualizes the intersections of four major development themes in the semantic wave: networking, content, services, and cognition.

R&D themes[2] include:

- *Networking* — Semantics to enable computers to configure and manage dynamic, persistent, virtual systems-of-systems across web, grid & P2P.
- *Content* — Semantics to make information interoperable, improve search, enable content discovery, access, and understanding across organization and system boundaries, and improve information lifecycle economics.
- *Services* — Semantics to enable computers to discover, compose, orchestrate, and manage services, and link information and applications in composite applications.
- *Cognition* — Semantics to make knowledge executable by computer; augment capabilities of knowledge workers; enable robust adaptive, autonomic, autonomous behaviors.

[2] R&D themes, including examples from North America, Europe and Asia, will be reviewed in *Semantic Wave 2006: Part-2.*

Figure-04 Semantic technology functions

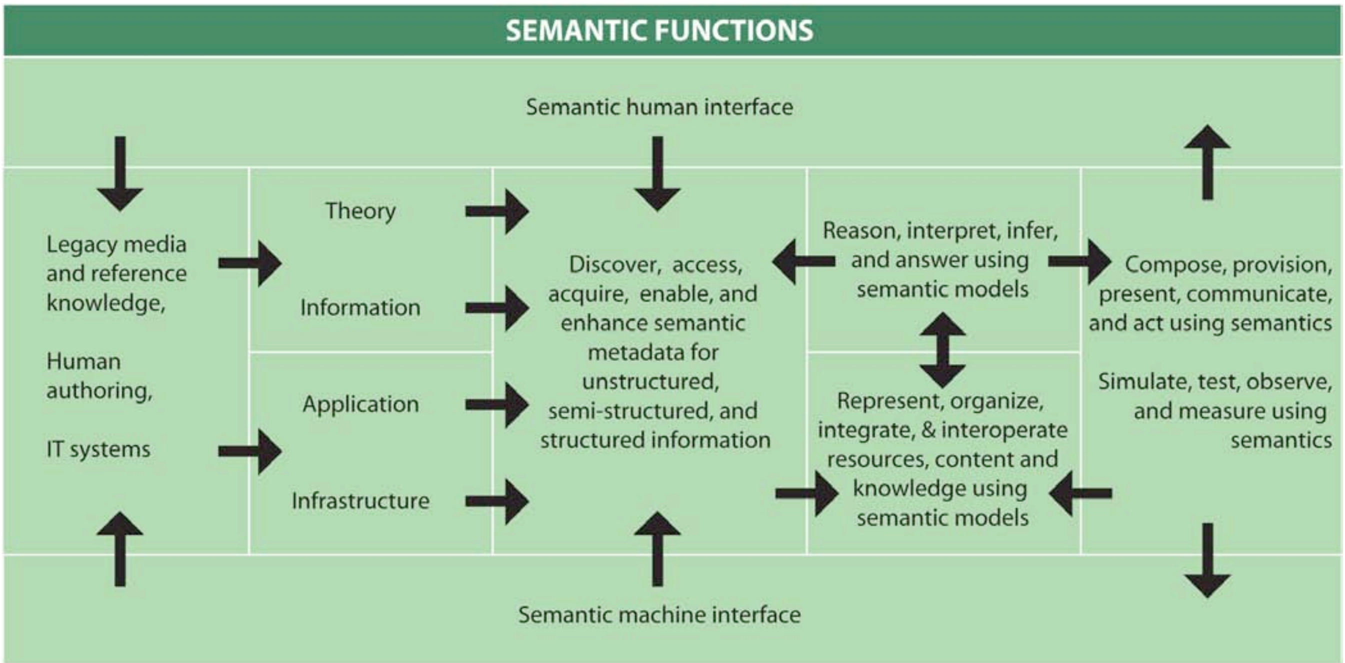

Source: MILLS•DAVIS

What are the functions of semantic technology?

Semantic technology functions are to create, discover, represent, organize, process, manage, reason with, present, share, and utilize meanings and knowledge in order to accomplish business, personal, and societal purposes.

Figure-04 highlights semantic technology functions:

- Semantic human interface at the top, and semantic machine interface at the bottom underscore the central goal of *making knowledge and meanings understandable and actionable by both humans and machines*.
- Capabilities to discover, extract, and model knowledge as well as enhance information with semantic metadata show from the left to the middle of the diagram. External sources include legacy media and reference knowledge, human authoring, and knowledge embedded in IT systems. Source-specific tools are used for recognizing patterns, syntax and structures within different data and language formats. Semantic tools provide capabilities for auto-recognition of topics and concepts, extraction of information and meaning, categorization, correlation, and mapping of interrelation between various sources of knowledge.
- The middle to right of the diagram depicts the capabilities to reason, interpret, infer, and answer based on using semantic models. Functions to compose, provision, communicate, and act based on semantics show to the right.

- Semantic technologies represent, organize, integrate and interoperate resources, content and knowledge. Organization of meanings makes use of taxonomies, ontologies and knowledgebases. These are relatively easy to modify for new concepts, relationships, properties, constraints and instances. Because semantic technologies integrate data, content, applications, and processes via a shared ontology, this minimizes development and maintenance costs.
- By using ontologies, semantic technologies can auto-discover and provision web services and functionality. Ontologies can link applications into composites that deliver a comprehensive view of situations with all data and information in context. Also, by representing meanings in language and media neutral forms, semantic technologies can auto-generate text, graphics, drawings, documents, and natural language dialogs. Similarly, they can auto-personalize, customize, and generate multiple versions of communications from the same knowledgebase.
- Semantic technologies reason via associations, logic, constraints, rules, conditions, and axioms that are represented in the ontology. This declarative structure allows reasoning in multiple directions. For example the same knowledgebase can be used to answer questions about how, why, and what-if as well as give factual responses. Given a question, semantic technologies can directly search topics, concepts, associations that span a vast number of sources, delivering results that are more relevant and comprehensive than searching with linguistic and statistical methods. They score higher in "recall" and "precision" measurements. Further, semantic technologies can deliver intelligence and answers to questions, not just provide lists of sources.
- A current direction is development of semantic functionality that can "learn" (infer and create new knowledge), simulate and test, and adapt behavior based on experience.

Figure-05 Semantic technologies impact all layers of the ICT stack

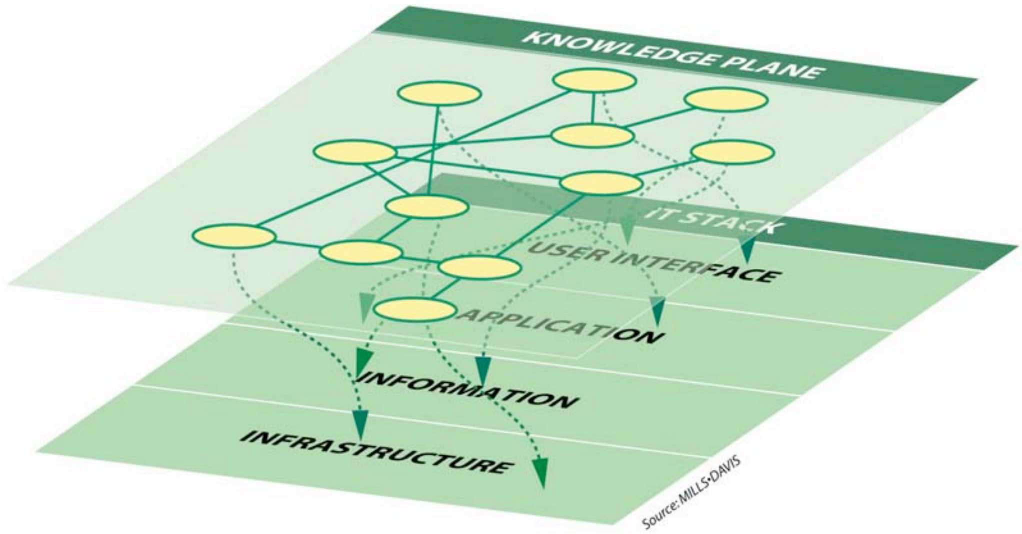

How do semantic technologies impact information technologies?

The first and most important thing to note is that semantic technologies do not just impact one layer or one part of the IT stack. They affect every aspect.

Figure-05 shows how semantic technologies impact all layers of the ICT stack. The knowledge plane (shown above) is a run-time semantic model, or web of knowledge about infrastructure, information, application process, user interface, system behavior, and other domains (shown below). The knowledge plane connects resources in and across each layer of the stack. It can be queried, interpreted and reasoned over both by people and machines.

This knowledge plane represents a new (knowledge-centric) computing paradigm. Orthogonal to the ICT stack, it opens a different dimension for architecture and development allowing cost-effective, sustainable solutions to problems of scale, complexity, connectedness, mobility, context, security, and interoperability.

4. Business Value of Semantic Technologies

Figure-06

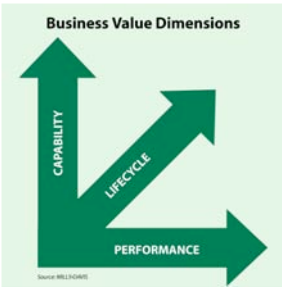

Business value dimensions

What is the business value of semantic technologies? Figure-06 depicts the framework for assessing the business value semantic technologies. It has three dimensions or axes:

- *Capabilities* enabled by semantic technologies and new solution patterns. New capabilities are the main value driver.
- *Lifecycle* economics of semantic solutions measured as the ratio of benefits to cost and risk — The lifecycle perspective focuses on development risk.
- *Performance* of semantic solutions measured by improvements in efficiency, effectiveness, or strategic edge. Performance focuses on returns.

In the following pages we discuss each dimension of business value in more detail.

Figure-07 Semantic capabilities drive business value

CHALLENGES	SEMANTIC CAPABILITIES	VALUE DRIVERS
Development: Complexity, labor-intensity, solution time, cost, risk	Semantic automation of "business need-to-capability-to-simulate-to-test-to-deploy-to-execute" development paradigm	Semantic wave intellectual property is about knowledgebases, less about code.
Infrastructure: Net-centricity, scalability; resource, device, system, information source, communication intensity	Semantic enablement and orchestration of transport, storage, and computing resources; IPv6, SOA, WS, BPM, EAI, EII, Grid, P2P, security, mobility, system-of-systems	In the semantic wave, infrastructure scale, complexity, and security become unmanageable without semantic solutions.
Information: Semantic interoperability of information formats, sources, processes, and standards; search relevance, use context	Composite applications (information & applications in context powered by semantic models), semantic search, semantic collaboration, semantic portals	Semantic interoperability, semantic search, semantic collaboration, and composite applications become "killer apps."
Knowledge: Knowledge automation, complex reasoning, knowledge commerce	Executable domain knowlege-enabled authoring, research, simulation, science, design, logistics, engineering, virtual manufacturing, policy and decision support	Executable knowledge assets enable new concepts of operation, super-productive knowledge work, enterprise knowledge superiority, and new intellectual property.
Behavior: Systems that know what they're doing	Robust adaptive, autonomic, autonomous system behaviors, cognitive agents, robots, games, devices, and systems that know, learn, and reason as humans do	Semantic wave systems learn and reason as humans do, using large knowledgebases, and reasoning with uncertainty and values as well as logic.

Source: MILLS•DAVIS

What capabilities of semantic technology drive business value?

Semantic technologies drive business value by providing superior capabilities (increased capacity to perform) in five critical areas:

- *Development* — Semantic automation of the "business-need-to-capability-to-simulate-to-test-to-deploy-to-execute" development paradigm solves problems of complexity, labor-intensivity, time-to-solution, cost, and development risk.
- *Infrastructure* — Semantic enablement and orchestration of core resources for transport, storage, and computing helps solve problems of infrastructure scale, complexity, and security.
- *Information* — Semantic interoperability of information and applications in context, powered by semantic models makes "killer apps" of semantic search, semantic collaboration, semantic portals and composite applications.
- *Knowledge* — Knowledge work automation and knowledge worker augmentation based on executable knowledge assets enable new concepts of operation, super-productive knowlege work, enterprise knowledge-superiority, and new forms of intellectual property.
- *Behavior* — Systems that learn and reason as humans do, using large knowledgebases, and reasoning with uncertainty and values as well as logic enable new categories of hi-value product, service, and process.

Figure-08 Semantics for development and IT governance

How do semantic capabilities impact development?

Semantic technologies fundamentally change the paradigm, tooling, practices, skills, and economics of solution development.

Figure-08 depicts a lifecycle methodology for semantic solution envisioning. The semantic approach is different. It is knowledge-centric rather than document-driven. Solution patterns based on semantic models (ontology) drive the process. The lifecycle development model is tuned to leaner, faster, build cycles, and a different mix of skills stressing declarative modeling over algorithmic programming.

Semantic lifecycle development is geared to deal with business needs and capabilities as the core discovery process, and to leverage semantic technologies and composite applications as means to deliver superior value. Unlike most modeling exercises, semantic development produces executable architecture, enabling rapid, iterative development. Also, design by example. It is not unusual to model (and deliver) working enterprise solution prototypes as part of the discovery process (e.g. 30 days), and production deployable solutions within 90-120 days. Semantic development is a fast, incremental, iterative approach.

Business capability exploration focuses on reaching agreement about basic concepts and terms that different groups use. As a vehicle for reaching agreement between stakeholders, an ontology supports multiple points of view as well as different vocabularies. Semantic models are inherently multi-perspectival and can generate controlled vocabularies and taxonomies as needed by different business lines, functional units, or communities of practice within the enterprise as well as across the supply chain.

Semantic development applies iteratively and incrementally, which accelerates return on investment. There is no need to "boil the ocean." Development is fast, incremental, iterative and non-invasive. Changes in ontology, data, rules, workflow, etc, can automatically update system functionality, without impacting underlying legacy systems and databases. This decreases time, cost, and risk to deploy.

The semantic development environment employs different tooling to produce different artifacts from different building blocks. A key advantage of semantic modeling is that it is immediately executable without writing code. Requirements, entities, relationships, capabilities, functions, processes, events, information, and business rules for composing an application are defined in a model that is directly executable by a semantic engine.

Building an "intelligence layer" allows delivery of capabilities and business value to business users by building composite applications over existing transaction silos. This preserves value of legacy investments. Favorite applications and tools are encapsulated and exposed in a composite application UI. Federated connectivity happens through standards-based database calls, APIs to apps, middleware interfaces, web services calls, etc. In short, with whatever IT environment the enterprise has. The knowledge plane models the essential business context, integration, relationships and business rules between applications, databases, and processes. Applications and data sources link to and interact with each other in real time and in context through the business ontology layer.

Dynamic semantic models can be reasoned over. Connections can be inferred. Also, such ontologies can be consulted by different applications at execution time, not just design time, make ongoing integration costs more linear rather than exponential.

Semantic models allow for efficient change management. Policies and general business strategy are tied to business ontology ensuring that systems support business objectives. Versioning of business ontology, data, rules, etc. can provide business with a snapshot of the business as of any point in time. Semantic modeling allows tracing system evolution. Metadata provides line-of-sight to metrics for analyses and reporting.

Semantic approach allows easier and more efficient training, maintenance, and support.

Semantic applications can be self-documenting and self-explaining. Documentation of the evolving system can be generated directly from the model, ensuring that as-is, and to-be documentation remains in synch, since both are driven from this evolving knowledgebase.

Learning to use the system is easier. In composite applications driven by semantic models, any object can be queried. All queries are contextual, derived from the semantic model (business ontology), and built in. Users do not have to learn separate queries for underlying systems that have been exposed through the semantic composite application. So using the system is easier. Semantic modeling empowers both self-help and superior service at the help desk, making it possible to deliver a higher level of support with significantly less resources.

Semantic development enables solutions with new capabilities, such as:

- Virtual infrastructure, semantically modeled middleware, and net-centric services and operations that reduce integration costs.
- Linking multiple structured, semi-structured and unstructured information sources through an ontology that allows users to search and access any source using their own business vocabulary, not the idiosyncrasies of querying individual systems.
- Non-intrusive, real-time integration and system-of-systems interoperability (internally, across supply chains) to provide advanced capability without requiring customers to rip-and-replace systems of record.
- Composite applications that enable knowledge workers to connect the dots, put information in context, gain 360-degree view of data, interact with information and applications in the context of their business process.
- Business aligned, rapid tactical implementation of strategic capabilities such as: enterprise IT integration, consolidation, and modernization; knowledge-centered customer-facing process; business intelligence, line-of-sight analytics, dashboards; exception management, case management; and compliance, cross-selling, command and control.

Figure-09 Semantics for infrastructure

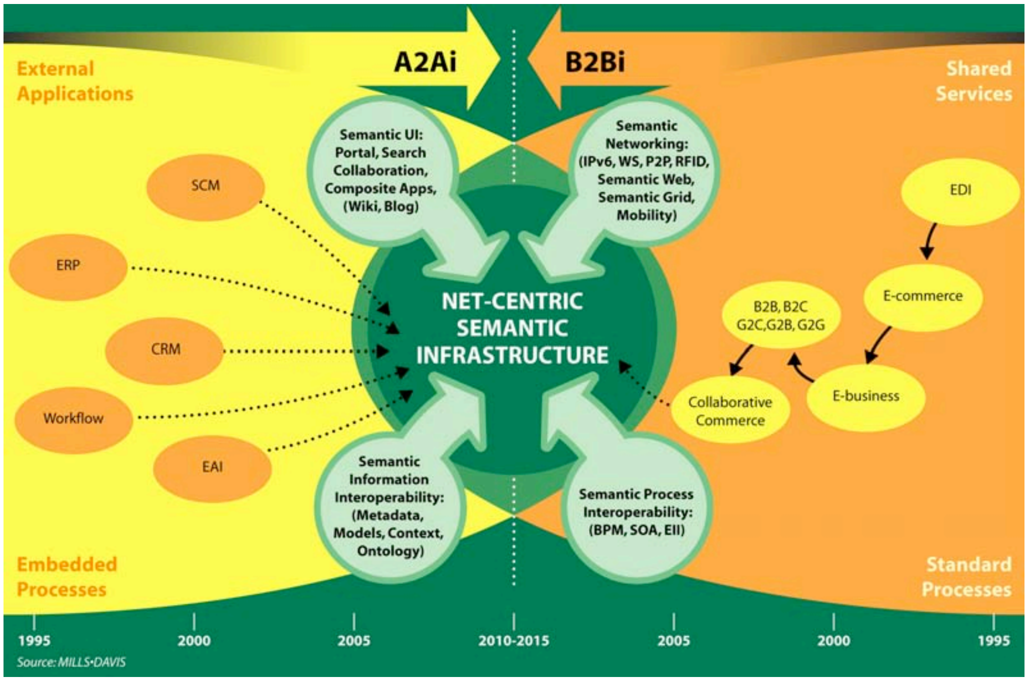

How do semantic technologies impact infrastructure?

Semantic technologies provide the glue to make net-centric operations practical.

Figure-09 depicts the evolution of application-to-application integration (A2Ai) and business-to-business integration (B2Bi) over the past decade and the current convergence of four streams of innovation towards net-centric semantic infrastructure.

- A2Ai began with the development of database management system technology as a means to share data and involves a number of paralleling developments in enterprise applications such as enterprise resource planning (ERP), workflow, customer relationship management (CRM), and supply chain management (SCM) systems. Efforts to integrate these began with proprietary solutions that could be enforced within the same hardware and operating system environment. As networking and client-server architecture grew in importance, vendor middleware and enterprise application integration solutions appeared.
- B2Bi has its roots in electronic data interchange (EDI) of the 1980s, but with the advent of the world-wide web (WWW) in the 1990s we witnessed a series of developments from e-commerce, to e-business, to B2B, to collaborative commerce.
- In the early 2000s, standards-based integration platform initiatives emerged from several quarters. The four streams depicted in this diagram have had differing objectives. *Network* initiatives focused on extending the Internet (IP), web services (WS), grid computing, and more recently by RFID, mobility, and Ipv6. *Process* integration initiatives focused on abstracting business process management (BPM) from applications and on service-oriented architecture (SOA). Information inte-

gration efforts focused on XML formats, meta-data standardization and enterprise information integration (EII). Portal integration initiatives focused on smart-client technology, integrated user interfaces (UI), and enterprise portal management (EPM).

- Today, it has become evident that these technology initiatives, by themselves don't pack enough punch to deliver all the capability breakthroughs and economic improvements needed for net-centric computing. These approaches lack a means to represent the semantics that are needed for machine understanding as well as for integration with other approaches. Whether the priority is mounting cost and complexity of maintaining current IT systems, development of new capabilities needed to meet challenges of network era computing, or strategies for achieving knowledge superiority, the combination of syntactical, structural *and semantic* interoperability is now required. Stated another way, if you want to connect information and processes and make them interoperable, first integrate the knowledge about them.
- The center of the diagram depicts where we are headed. We are moving from legacy stovepipes and proprietary services to net-centric semantic infrastructure. The four streams of ICT innovation are converging and semantic technologies are becoming a core aspect of each of them. The web is growing a semantic web. Grid computing is becoming semantic grid. The same with P2P, RFID, mobile computing, and IPv6. Semantics are now part of the roadmap for SOA, BPM, and EII. Semantics are the path forward to information interoperability. Semantic technologies are transforming the user interface, portals, search, collaboration (e.g. semantic Wiki) and enabling composite applications linking diverse applications and data sources.

Semantics for enterprise architecture

Semantic technologies transform enterprise architecture (EA) from a reference document to an executable intellectual backplane for IT governance.

Gartner defines an enterprise architecture tools as having:

- A repository in which to store information about the business, applications, data and technologies
- A metamodel to structure this information
- The ability to represent information in the repository in graphical and textual forms

The purpose of enterprise architecture is to comprehensively define the mission, organization, functions, performance, processes, information, technologies, and social aspects of the entities that IT supports.

However, the millions spent to date to develop enterprise architectures as a basis for IT modernization have largely resulted in manual compliance exercises, producing reference documentation, disconnected from operations and management systems, and delivering no capability to business users.

The semantic EA approach is different. Its tools produce semantic models that are executable, operational, and integrate directly with existing applications and information sources. Semantic models connect systems and information sources, and provide a knowledgebase that enables line-of-site analysis and management of the transition from as-is to to-be stages of the enterprise IT environment.

The semantic EA approach also includes semantic auto-discovery and mapping of legacy IT artifacts and documentation. This gives visibility and eliminates 1/5 to 1/3 of cost of as-is modeling, compliance auditing, and steady-state maintenance projects. Semantic discovery applied to IT artifacts is the capability to scan source libraries, data schemas, and documentation, comments, etc. in order to identify unique artifacts, link and map dependencies, and do latent semantic indexing of the "as-is" world. The result is a repository of metadata (RDF/OWL), a very flat ontology that enables semantic (concept) search using business terms, without having to know the (often cryptic) as-built naming established by programmers. This could be thought of as a sort of "Google for IT" process that works bottom up, and also allows mapping linkages to enterprise architectures, or other governing models.

Executable enterprise architecture provides operational enterprise integration and semantic interoperability across mission, function, process, information, and resource aspects of enterprise management. It delivers an integrated, semantic model-based operational capability for cross-enterprise investment planning and control, IT portfolio management, and compliance. This provides huge leverage on application management and modernization costs. Semantic EA establishes an executable enterprise knowledgebase that enables line-of-sight analyses, analytics, and automating of alignment and compliance with enterprise goals and policies, saving labor, cost and time.

Semantics for integrating systems and data

The low-hanging fruit over the near-term is semantic interoperability of information and processes. Semantic integration changes the economics of bridging disparate data sources and legacy systems.

First, instead of point-to-point integration that results in N^2 complexity growth as the number of components increases, linking through a knowledge layer creates a hub-and-spoke model that reduces the number of new interfaces to a linear progression.

Second, much of today’s motivation for using web services is to take cost, time, and effort out of the processes of integrating applications. However, it still requires people to research and construct the interfaces, even using directories to locate services. Semantic web services enhance the value proposition. Service ontologies provide a way to take people out of the loop of configuring and managing the integration, resulting in reduced total cost of ownership (TCO).

Figure-10 Semantics for information, knowledge, and reasoning

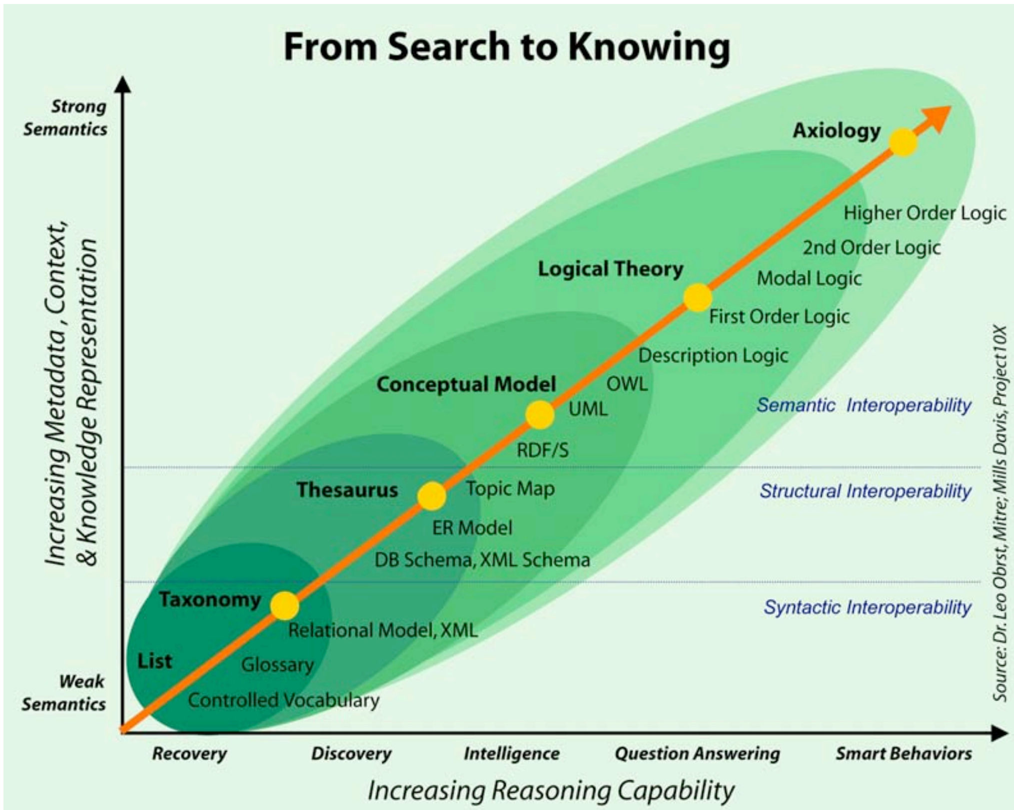

How do semantic technologies impact information and knowledge?

Semantic capabilities enhance value and improve the lifecycle economics of information and knowledge.

Semantic enablement of information can enhance authoring, search, discovery, access (or sharing), aggregation, understanding, and communication of information. It imparts new capabilities for knowledge work automation and knowledge worker augmentation.

Figure-10 depicts a spectrum of interoperability and reasoning capabilities from search to knowing:

- From bottom-to-top, the amount, kinds, and complexity of metadata, modeling, context, and knowledge representation increases.
- From left-to-right, reasoning capabilities advance from (a) information recovery based on linguistic and statistical methods, to (b) discovery of unexpected relevant information and associations through mining, to (c) intelligence based on correlation of data sources, connecting the dots, and putting information into context; to (d) question answering ranging from simple factoids to complex decision-support, and (e) smart behaviors including robust adaptive and autonomous action.
- Moving from lower right to upper left, the diagram depicts a spectrum of progressively more capable categories of knowledge representation together with standards and formalisms used to express metadata, associations, models, contexts, and modes of reasoning.

- More expressive forms of metadata and semantic modeling encompass simpler forms, and extend their capabilities. For example, the semantic web standard OWL encompasses the ability represent glossaries, taxonomies, thesauri, subject ontologies, and the semantics of specific XML schemas, database schemas, entity-relationship and UML models. On the other hand the OWL standard is not capable of representing the full spectrum of logical theory, nor higher order modes of reasoning. But, other formalisms exist that can express these capabilities.
- As the amount and expressive power of the semantics and knowledge increases, so does the value of the reasoning capacity it enables.

As information volume explodes and lifecycle costs surge, businesses face crucial issues, for example, how to:

- Mobilize and make sense of distributed content assets from diverse provenance
- Ensure content is trusted, and authoritative for its intended use
- Improve content accessibility and search quality
- Leverage knowledge worker content use in specific contexts
- Improve the economics authoring, provisioning, and distribution costs
- Achieve content interoperability and integration across diverse sources that have been classified and indexed by different communities.
- Reason over content and knowledge sources to answer questions and support decision making
- Cope with information overload
- Achieve multiple returns on content investments

Metadata

Content is anything that is written, depicted, filmed, recorded, animated and stored in some media. Digital content is any written, depicted, etc. content whose physical properties can be substituted by computer-processable descriptions. E.g. a digital recording on a CD is a binary description of the sound.

When digital content's primary information (e.g. the music, video, text, etc.) is enhanced by secondary information about the content, this secondary information is called metadata.

Metadata is not just one thing. Various kinds of metadata enable differing capabilities, for example:

- Domain knowledge provides a framework for interpreting the meaning of the content from different perspectives.
- Modelling the context of use enables software to organize content to task, interest, or preference.

- Media resource knowledge enables packaging content for presentation across different media.
- Behavioural knowledge allows the computer to sequence communications and manage dialogs.
- Provenance and rights metadata is key to establishing trust and maintaining security, and enabling commerce.

Semantic discovery

Semantic discovery applied to content includes the capability to:

- Scan all types of data and information formats encountered on the desktop or Intranet, or web, and to extract concepts, relationships, and constraints; create metadata repositories; and represent (or model) knowledge in the form of ontologies.
- Ingest, interpret, and merge metadata and ontologies from internal and external sources -- thus providing a basis for (machine understandable) knowledge sharing, and commerce in knowledge assets.
- Enhance content with semantic (metadata) through inference and linking facts and concepts with domain ontologies and knowledgebases.
- Identify, extract, and model mediating structures, business rules, axioms, and constraints, from repositories, as-built applications and data structures, and/or authoritative knowledge-rich documents such as policy, regulation, or law, standards, and reference knowledge.

Information intelligence

Semantic capabilities enable information intelligence (information in context of need) through aggregation, integration, and interpretation of diverse data sources. The spectrum of requirements includes:

- *Sense-making* — Extract knowledge and tag metadata based on statistical, language-based, semantic, and knowledge-centered approaches. Enable sharing and interoperability at this level through data services that parse formats, match patterns, distinguish features (such as parts of speech), apply linguistic and statistical methods, etc. Semantic adaptation services mine and extract knowledge and semantics from data sources, or otherwise add semantic metadata of various kinds to the data. Semantic integration services link information, metadata, and semantic models.
- *Information sources* — discovery, access, and understanding of structured, semi-structured, unstructured information sources. Sources are federated and distributed.
- *Information structure levels* — Signal, data, content, metadata, model, and semantic model; sharing and interoperability span a continuum of contexts.

- *Search contexts* — Semantic query services access, navigate, and reason over semantically enabled content to be provisioned to various client applications. Retrieval, discovery, intelligence, question-answering, and decision-support reasoning, and thus a need to enable exploitation of content interoperability at increasing cognitive depths (Significant literature exists, including the 12 years of testing conducted through NIST's TREC conferences, that documents the efficacy of different approaches to increasing precision and recall, relevance, speed, and scalability.)
- *Sharing contexts* — encompasses: (a) general search, (b) task or context-based search and line of thought navigation, (c) composite applications providing 360-degree view, integration of structured and unstructured information in context of need, and interaction with information in user-determined context involving processes, tracking; and (d) mission and time-critical situation awareness, reasoning and trade-off assessments, and decision-support, and (e) autonomic, adaptive, and autonomous system behavior.

Universal knowledge technology

Over the next decade, we can expect rapid progress towards a *universal knowledge technology* that can provide a full spectrum of information, metadata, semantic modeling, and advanced reasoning capabilities.

Why? First, there exist significant unmet needs. Very large-scale knowledgebases, complex forms of situation assessment, sophisticated reasoning with uncertainty and values, and autonomic and autonomous system behavior pose challenges that exceed the capabilities and performance capacity of current open standards approaches. Second, no good reason exists for settling for only a portion of the capability spectrum when we can just as easily have the whole thing.

Universal knowledge technology will be based on a physical theory of knowledge that holds that knowledge is anything that decreases uncertainty.

The formula is: *Knowledge = Theory + Information*.

- *Theories* are the conditional constraints that give meaning to concepts, ideas and thought patterns. Theory asserts answers to "how", "why" and "what if" questions. For humans, *Theory* is learned through enculturation, education and life experience and represents 85% of knowledge content.
- *Information,* or data, provides situation awareness — who, what, when, where and how-much facts of situations and circumstances. Information represents only 15% of knowledge and requires *theory* to define its meaning and purpose.

What distinguishes universal knowledge technology is that it enables both machines and humans to understand and reason with any form of knowledge, of any degree of complexity, at any scale.

Universal knowledge technologies will capture meanings and knowledge separately from natural language, media, data forms, or computer program code. Ideas are represented precisely using nth-order conceptual notation with guaranteed minimum bit-level encoding approaching Shannon limits. Unlike relational databases or object databases, resource requirements increase linearly with knowledge quantity and complexity, with no combinatorial explosion or performance degradation as the knowledgebase scales.

Using nth-order conceptual notation rather than name spaces allows the syndetic precision needed to overcome inherent ambiguity and elasticity of natural language. Also, it provides the meta-level modeling capacity needed to semantically integrate a myriad of purpose-built formalisms. This allows universal knowledge technology to encompass and to be compatible with a broad spectrum of open standards.

What does this mean? Universal knowledge technology will embrace and extend capabilities of both information technology and the semantic web. The capacity to assimilate all knowledge domains, all kinds of theory, and all types of information allows universal knowledge technology to excel at infrastructure-, information-, knowledge-, and behavior-intensive tasks. It enables solutions to interoperability, integration, and infrastructure proliferation dilemmas of current information technologies. It overcomes the limitations to scalability, complexity, and reasoning inherent in current semantic web standards.

Universal knowledge technology will reason declaratively using all forms of induction, abduction, and deduction incorporating logic, uncertainty, conflict, and values. It can apply ontologies passively to enhance information search, data exposure, or social networking. And it can execute knowledge actively for research, robust simulation and decision-making in sciences, engineering, professions, governance, management, and value-based pursuits of all kinds.

Universal knowledge technology can capture, fully integrate, and reason with any concept, thought pattern, theory, meaning, and datum found in books, documents and databases — as well as the very "how we do it here" knowledge in the minds of professionals, researchers, employees, and communities of interest. This brings the capacity to:

- Assimilate an entire corpus of theory being developed through community-based scientific research, providing a basis for in-silico experimentation, simulation and testing.
- Capture complex system of system engineering knowledebases from documents, databases, and team member contributions, providing a foundation for design-advisors and simulation-based virtual manufacturing, and command and control.
- Capture reference knowledge to enable, for example, legally-defensible diagnoses and reasoning in the professions, or in-depth decision-support to examine complex trade-offs and unexpected consequences for managerial and policy-level decision-making.

Figure-11 Semantics for information-intensive knowledge work

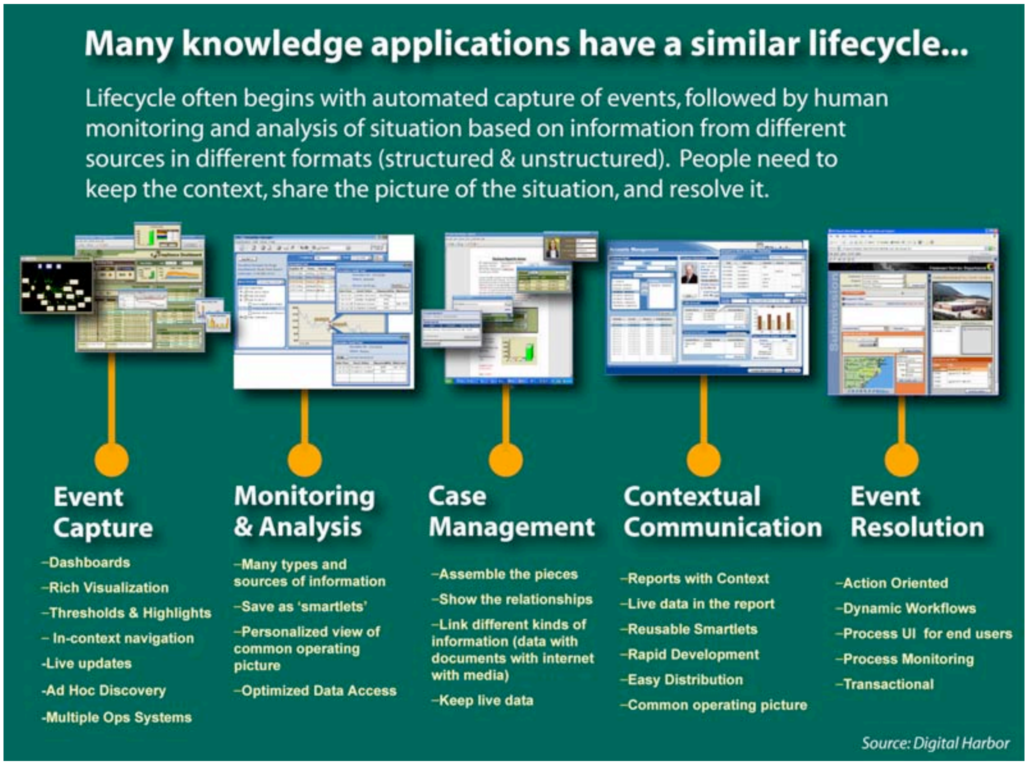

How do semantic technologies impact information-intensive knowledge work applications?

Semantic technologies deliver a "killer applications" for the semantic web—collaborative composite applications, driven by semantic models, enable knowledge work automation and knowledge worker augmentation.

Figure-11 depicts the lifecycle of information-intensive knowledge work applications. The lifecycle often begins with automation capture of events, followed by human monitoring and analysis of a situation based on information from different sources in different formats (structured and unstructured). People need to keep the context, share the picture of the situation, and resolve it.

Composite applications fuse data and services from multiple applications, correlate information in context, drill down and across in real-time, ask questions across databases, and infer links across systems. Knowledge workers see information in context. They have live application functionality in the UI, giving them real-time interaction among systems, They can ask questions on any object in language they understand giving them a 360-degree view and the ability to understand and act in context.

Composite applications deliver capability and business value directly to the business user. Knowledge workers do not care about IT architecture. They want capabilities. The capabilities they need demand information integration, system-of-system interoperability and power-to-the-edge. There is no way they can define all requirements in advance. That's why businesses need a way to deliver benefits directly to end-users that is fast, affordable, incremental, and non-invasive.

Figure-12 Composite applications powered by semantic models

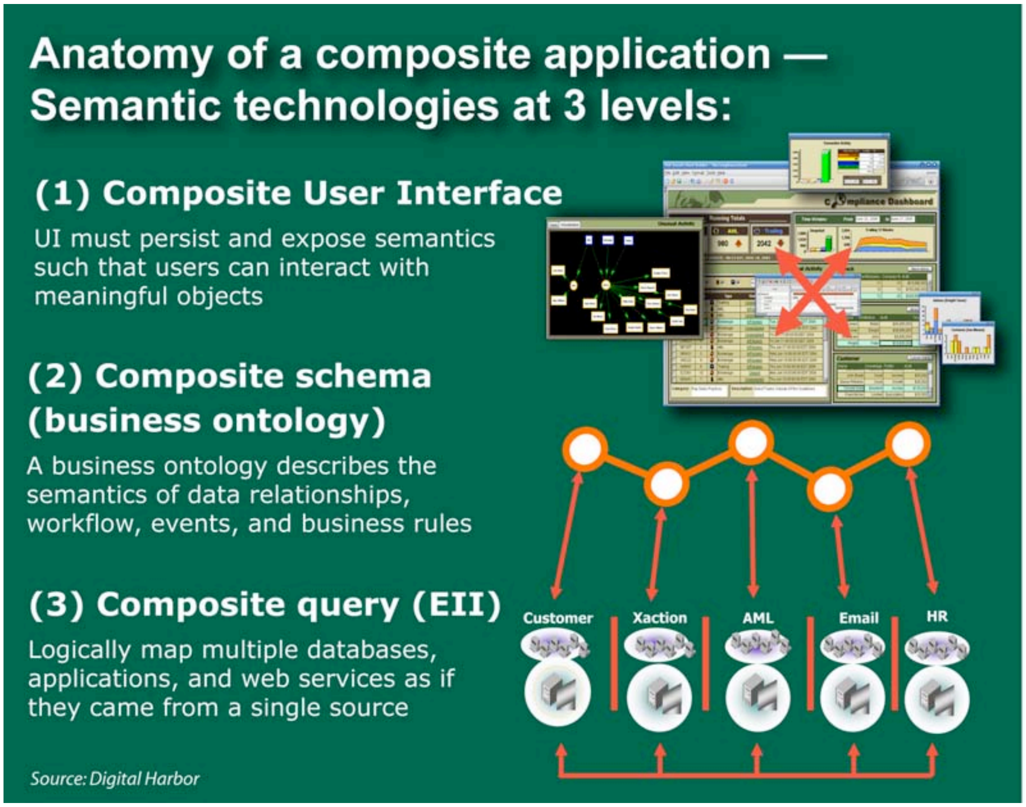

Figure-12 depicts how composite applications incorporate and utilize semantic models (i.e. knowledge plane). They do so at three levels:

- User interface persists and exposes semantics such that users can interact with meaningful objects;
- Business ontology describes the semantics of data relationships, workflow, events, and business rules. It leaves data in its physical source(s), but logically relates different kinds of information.
- Composite queries logically map multiple databases, applications, and web services as if they came from a single source.

Building composite applications with semantic models is 5X faster to develop than contemporary methods. Teams are lean. There's reduced project risk. Having a tactical, non-invasive, iterative solution for strategic modernization empowers IT. Composite applications link information in context, empowering new categories of knowledge worker capabilities for: exception handling, emergency response, compliance, risk management, situation assessment, command and control.

Semantics for collaboration

Semantic Wikis illustrate trends towards semantic-enabled collaboration and information sharing.

A Wiki (from WikiWiki, meaning 'fast' in Hawaiian) is a set of linked web pages, created through the incremental development by a group of collaborating users, as well as the software used to manage the set of web pages. A Wiki:

- Enables web documents to be authored collectively
- Uses a simple markup scheme
- Does not publish content instantly, once an author submits a page to the Wiki engine
- Creates new web pages when users create hyperlinks that point nowhere.

A semantic Wiki creates an overlay network structure (aka knowledge plane) that defines concepts, attributes, and relationship of the underlying content in the Wiki. Relationships become explicit as links. When navigating the semantic layer, links followed depend on the task and may change over time. Also, criteria to follow a link depend on the Wiki specifications in the overlay layer, which users that manage the overlay layer can modify as needed.

Semantic Wikis enhance collaboration and information sharing by providing capabilities such as:

- Concept-based rather than language-based searching: queries span vocabularies, languages, and search engines
- Question answering rather than simple retrieval. Also, overlay ontologies and knowledgebases can integrate with major web searching engines
- More richly structured content navigation, including multiple perspectives, multiple levels of abstraction, dependency/contingency relationships, etc.
- Easy visualization of content structure (categories, taxonomies, semantic nets, etc.). Direct editing of content structure.
- Mining of semantic relationships in content.
- Wiki content linked to dynamic models, simulations, visualizations.
- Wiki content linked with external repositories, file systems (e.g. personal desktop, enterprise servers, web sources, semantic-enabled feeds [e.g. RSS])
- Richer user access/rights models, including reputation systems.

Semantics for knowledge-intensive applications

How do semantic technologies impact knowledge-intensive applications?

Knowledge applications are the forward edge of the semantic wave. Semantic capabilities enable new categories of knowledge work automation and knowledge worker augmentation. Knowledge computing generates benefits up to an order of magnitude greater than those provided through semantic interoperability of systems and information sources, including:

- New categories of tools, applications and processes based on knowledge computing
- New categories of intellectual property — computable knowledge assets
- Quantum shifts in the capabilities, competitiveness, and economics of labor

Figure-13 illustrates navigation of semantic patterns of thought across federated knowledgebases. The problem shown is from a threat neutralization scenario. The left-hand column overviews the concept of semantic browsing. The central panel displays a color-coded honeycomb of integrated knowledgebases whose aggregate complexity is several orders of magnitude greater than a large enterprise database. Different colored hexagons represent separate projects that captured and progressively integrated each segment of reference knowledge. The overlay of numbers and connecting arrows depicts a rational path by which a question posed in the scenario was analyzed and both expected and unexpected trade-offs explored. This question-answering path also displays down the right-hand column and across the bottom section of the chart together with thumbnail screen shots of the browser window at each point in the analysis.

Figure-13 **Semantic question answering**

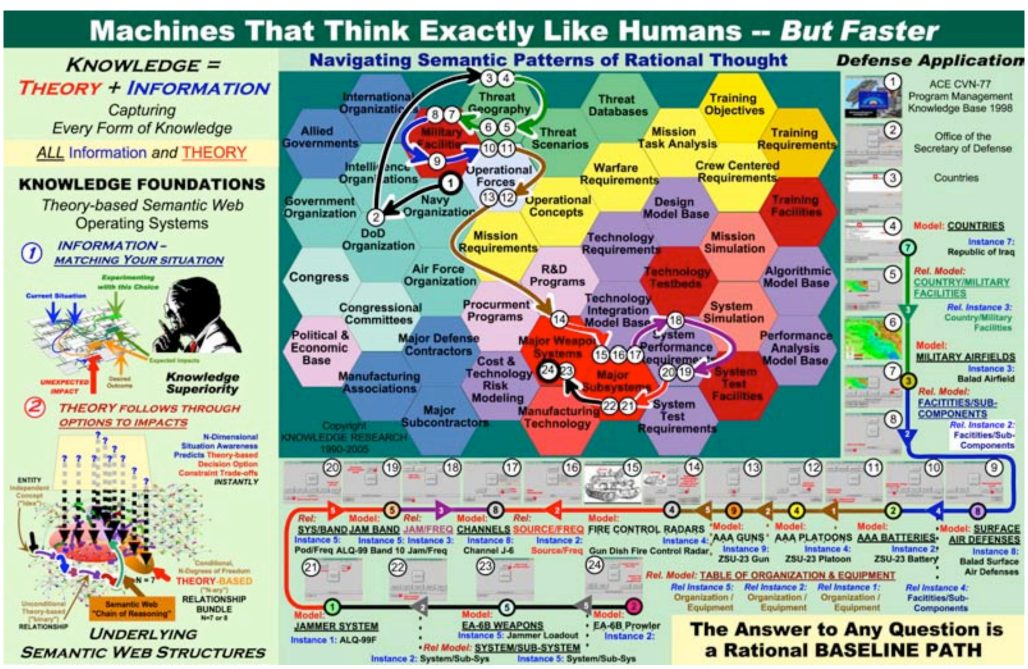

Knowledge work automation and knowledge worker augmentation

Semantic technologies and knowledge assets will transform the economics of labor, including the cost of education, personnel acquisition, productivity, and labor rates.

The spectrum of current strategies for improving workforce productivity and managing labor costs includes:

- Mechanization seeks to substitute capital investment in machines for labor.
- Outsourcing seeks to exploit differentials in labor rates and other costs among different geographies and business entities.
- Labor transitions (e.g., from professional to paraprofessionals in law and medicine) seek to substitute less-skilled workers for higher-cost workers in certain tasks.
- Service automation seeks to displace labor or maximize productivity.
- Self-service seeks to offload labor costs to the customer or supplier.
- Information technology seeks to improve labor productivity through digitization, automation, integration and optimization of information-based tasks and activities.
- Education, training and distance learning seek to transfer knowledge efficiently from sources to empower new generations (of labor).

Semantic technologies (i.e., executable knowledge embedded in tools, processes and infrastructure) will both accelerate and dramatically intensify the impact of all of these approaches for dealing with labor costs.

At the same time, semantic wave technologies will promote an unprecedented degree of career mobility and enhanced productivity at all levels of the job market. Given a professionally adept machine backup, early-career specialty training will be substantially shorter, but adaptive mid-career training will be constant. This is good news.

However, labor transitions will impact professions, management, and technical ranks—categories that previously have been less impacted than agriculture, manufacturing and service industries. Sustainable careers for highly educated, specialized professions will shift toward new knowledge discovery and marketable knowledge-asset creation.

Executable knowledge will become the basis for new categories of research, analysis, planning, design, diagnosis and decision-management tools. Knowledge tools have broad application. There are as many domains for their application as there are:

- Industry sectors and segments—government, manufacturing, services, energy, publishing, etc.
- Job categories—by role and responsibilities within an organization.
- Functions—such as decision-making, research, design, planning, analysis, marketing, sales, support.

- Disciplines—including management, projects, engineering, accounting, finance, software development, medicine, law, scholarship, etc.
- Hobbies and interests—gardening, home improvement, entertainment, games.

Semantic wave knowledge tools offer new capabilities. For example:

- *Law* — Having executable theory, legislation, and applicable case law in-the-computer could enable a legal researcher to both retrieve case law that is relevant to the brief *and* see its reasoning applied to the case at hand. During discovery, being able not only to identify rapidly all relevant legal documents, but also more thoroughly explore relevant approaches to arguing the brief, would both reduce risk and increase productivity of legal professionals. Also, one can imagine less-educated paraprofessionals with knowledge-based tools being able to perform diagnoses and other key functions of professionals, in legally defensible ways, resulting in cost savings.
- *Engineering* — Having theory-and-information-in-a-computer, leads to new kinds of design-build processes for manufacturing, capital projects, architecture, and engineering, resulting in a faster, more efficient lifecycle that could scale to handle very large complex projects.

  A key problem of current processes is that they are document-centric rather than knowledge-centric. To illustrate, across the engineering lifecycle, a part design can translate into hundreds of drawings, schematics, and documents prepared for different disciplines, or usages at different stages. Current document-centric workflow utilizes CAD and CAE tools as electronic pencils for creating and recreating documents. As project size and complexity grows, internal document maintenance and management consumes 80-90% of resources.

  In contrast, semantic wave knowledge tools capture, represent, and maintain total product knowledge in a language-neutral, federated repository. Semantic applications generate all categories of engineering drawings, specifications, project documents, and technical literature as needed. The result was an up to 5-10X faster design, build cycle with up to 5-10X reduction in project costs. Far fewer engineering resources were needed for projects of any given size. Knowledge-centered engineering enabled control of larger and more complex projects than with conventional methods. The ROI resulted from taking huge amounts of labor, cost, and time out of the process. Lifecycle knowledgebase removes errors and inconsistencies; gives visibility to all parts, versions, and phases of the project; and stops knowledge erosion due to personnel changes.
- *R&D* — Having theory-in-a-computer immediately calls for some way to test it. Simulation is the preeminent way to test, hence a low-hanging fruit for semantic wave knowledge tools. Semantics and the abundance of theory concerning every physical, rational and social process will make knowledge-based simulation a central subject not only in the sciences and engineering, but of every argument on plans, policies, strategies, new law, economics, social values, etc.

Knowledge computing

Key trends in the computation of knowledge include: (a) modeling all forms of knowledge, not just the relatively limited models needed to form enterprise architectures and composite applications, (b) zero code, declarative (non-algorithmic) application development, (c) tractable semi-automated and automated costs of knowledge acquisition, (d) high-performance reasoning over massive knowledgebases, (e) intelligent agents, and (f) adaptive, autonomic, and autonomous human and robotic systems.

These developments will impact capabilities for research, analysis, design, engineering, virtual manufacturing, logistics medicine, law, management, and advanced decision support to power business strategies based on knowledge superiority, performance augmentation, autonomous systems, and labor transitions.

For example, research in the life sciences is moving towards system of systems rather than a focus on individual parts. The process is an integrated cycle involving hypothesis, simulation, and observation.

"Pathway" is a common word in medicine and biology. It can be drawn as a picture to explain a complex process. It implies a cascade of events, where a "signal" is sent from one point in a cell/organism that creates a response. Pathways are hugely complex, traversing multiple disciplines, each with its own theories, terminology, and experimental data. Gaps in knowledge, uncertainty, and even conflicts between theories are a normal part of research.

Life sciences embrace ontologies as tools for e-science because they provide seamless formalism from the abstract to the concrete. The open world assumption embraces federated modeling, where models can be independent from solution, therefore can be upgraded, combined, evolved, reused. Ontologies provide direct mapping between concept and solutions, and offer constructs that can support context, which is important given the many different disciplines and perspectives involved.

Key challenges for knowledge computing for the life sciences include:

- Research data and knowledge is exploding faster than Moore's law.
- Even a small vertical slice of cell biology crosses many disciplines of knowledge.
- Knowledge / Data is distributed globally across thousands of databases, taxonomies and ontologies
- Observational science with large body of clinical and research knowledge trapped in millions of publications
- Private-public tension — Systems of private and public knowledge evolve rapidly, feed each other, and cannot be compromised

Figure-14 Executable knowledge powered by semantic agents

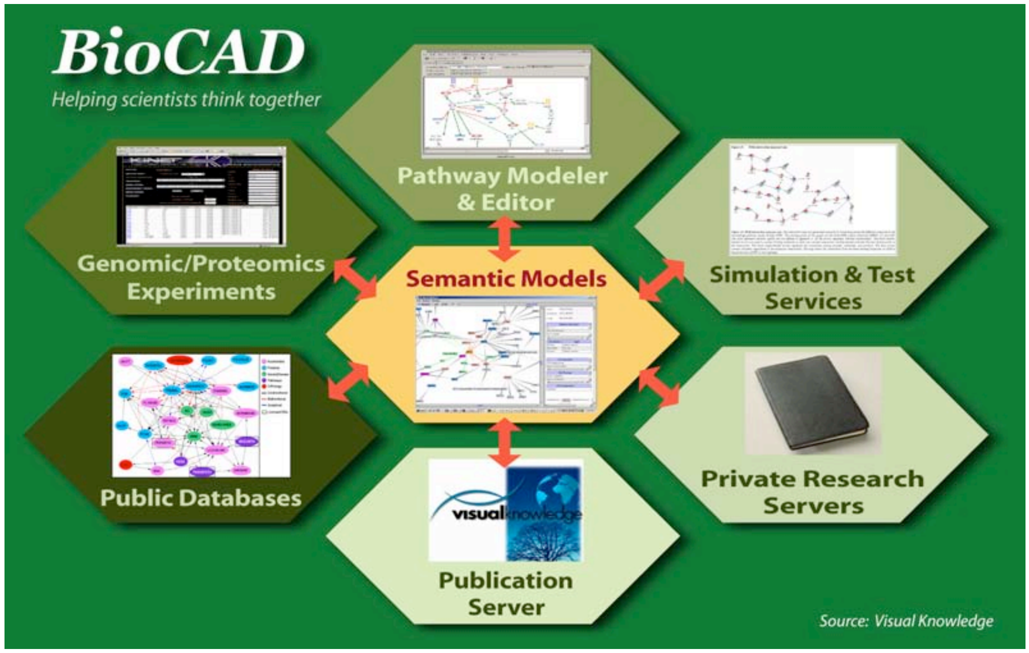

Figure-14 illustrates knowledge computing trends in the life sciences. The goal is to move to executable knowledge-based silicon laboratories where researchers actively create, experiment with, simulate, and test new hypotheses using computer-executable theories and experimental data contained in published literature, databases and private sources.

The BioCAD example depicts research knowledge computing using:

- *Public databases* — ontology-driven conversion integration of dozens of online community databases
- *Semantic models* — Real time inferencing for pathway analysis of cause and effect; currently, more than 27 million concepts hyper-integrated
- *Experimental data* — live production transactional system integrated data warehousing, mining, with real time supply chain OLTP on top of 27M concept ontology
- *Pathway modeler & editor* — ontologies are built by visual metaphor using easy to use drag / drop interface; there is international curation of resulting models; researchers can map experimental results onto visual pathways
- *Simulation and test services* — examine emerging research against what is already known about gene regulatory nets, protein signaling and metabolic pathways; chemical to cellular concepts
- *N-tier semantic publication facilities* — role/team based security, semantic change control, ability to develop/test/revise proprietary hypothesis; support for multiple experimental protocols including micro array, PCR, blots, mass spec; and publishing to private or public collaboration network

The BioCAD research environment illustrates a totally semantic application. Its DNA is declarative knowledge. There is zero program code. Everything happens through semantic agents powered by living ontologies. The semantic platform is able to handle federated development as well as thousands of concurrent users who actively contribute to the growing knowledgebase as a by-product of their normal research activities. Semantic technology automates comprehensive version and change management. The system deploys on PCs, Cray supercomputers, or research Grids exploiting massive parallelism. Knowledgebases and reasoning capabilities scale easily to terabytes.

Figure-15 Semantics for systems that know what they're doing

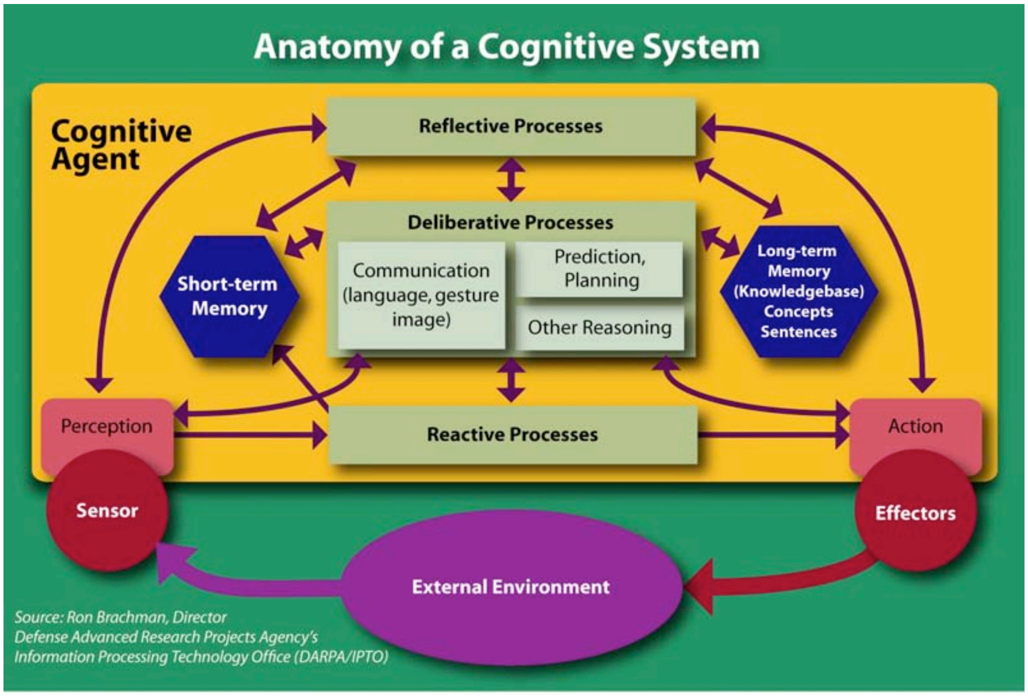

How do semantic technologies impact system behaviors?

Semantic capabilities enable new categories of "cognitive systems" to power robots, intelligent agents, and systems that know, learn, and reason as people do.

Figure-15 depicts the anatomy of a cognitive system. According to Ron Brachman, Director of the Defense Advanced Research Projects Agency's Information Processing Technology Office (DARPA/IPTO), cognitive systems are *systems that know what they're doing.*

A cognitive system is one that can:

- *Reason*, using substantial amounts of appropriately represented knowledge
- *Learn* from its experience so that it performs better tomorrow than it did today
- *Explain* itself and *be told* what to do
- Be *aware* of its own capabilities and *reflect* on its own behavior
- *Respond* robustly to surprise

Knowledge computing is about systems that know what they're doing, can reason as humans do, and can learn. The goal is the delivery of robust, adaptable, transparent, supervisable, autonomous intelligent systems with the ability to acquire, through experience, models of the world (including other entities and self), and use them productively to solve novel problems and deal successfully with unanticipated circumstances.

Acquiring models means learning by observation, exploration and experiment, teaching and coaching, or reading. Using models means reasoning including "mental simulation and testing," hypotheticals, plausible inference, logical thinking, and value-based trade-offs.

A key issue for developing cognitive systems is the distinction between architected or human modeled ontology on the one hand, and emergent, largely machine-automated knowledgebase construction, on the other.

We come from an era where databases and IT applications were always hand-built, and architected by humans. This is what most IT people know how to build. Also, our strategy for developing IT systems has been to build algorithms that we execute. The historical reasons for this were bandwidth and cost of computation. Major enterprise systems only consist of a few thousand relational tables.

In the net-centric, semantic era we move in a different direction — towards massive, declarative knowledge structures with billions and trillions of concept meshes. Here, everything known is always, already pre-computed; here situation awareness — delivers the facts and information which constrain outcome of questions posed including expected and unexpected outcomes; here there is no search — the answer to a question is a rational path over n-ary relationships.

Knowledge science shows that the solution to complexity exists, but cannot be achieved through layered abstraction. Rather, the solution is to capture all knowledge used by both humans and machines in n-dimensional theory-based declarative semantic form. This will to new software and computing architectures, while providing unlimited capacity to build massive stores of knowledge, of all kinds. Knowledge science tells us what knowledge is, how to represent it the most efficient way, how to reason with it in the fastest and most cost-effective way, and how to calculate and prove that no more resource-efficient way to compute knowledge exists.

Figure-16 Semantics for knowledge commerce

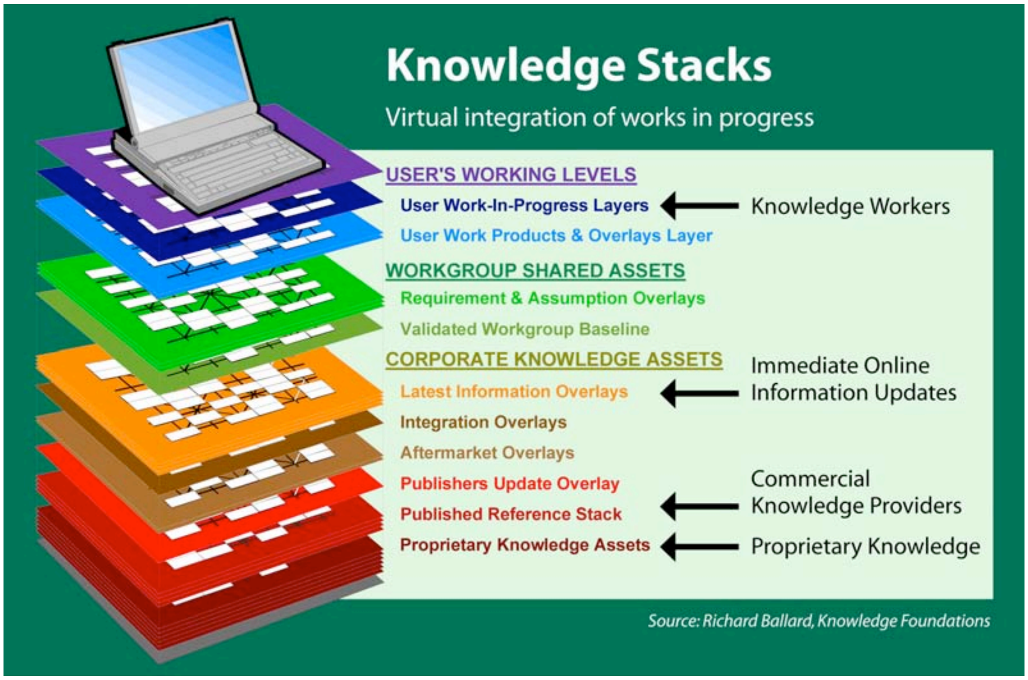

How do semantic technologies impact intellectual property?

Semantic technologies enable new categories of intellectual property as a foundation for knowledge sharing and knowledge commerce.

Knowledge that is understandable by humans and executable by machines represents a huge new market opportunity based on a new category of intellectual property — the knowledge asset. Declarative and computable, knowledge assets will become a key value driver for governments, communities of interest, enterprises, and individuals.

Putting both theory and information into a computer executable form creates a wholly new experience for the user. It's like the difference between reading a book about playing a game of chess and having an expert advisor to help you strategize and play the game better than ever before.

Computable knowledge can become an active (not passive) asset that is self-evolving and self-learning, and increases in value as it is used. Commercial, public, and private mechanisms will emerge to facilitate knowledge sharing and exchange.

Figure-16 shows a knowledge work environment wherein a collection of knowledge asset layers has been gathered, organized, integrated and securely worked upon by any number of contributing users. A session overlay creates this stack by dynamically linking together the working layers. Linkages are virtual, so they do not alter underlying copyrights, security, and permissions. In this example, user work-in-progress layers may be updated, while proprietary product layers are treated as "read only" within such stacks. Combined with encryption and digital rights management knowledge stacks provide an enabling infrastructure for knowledge commerce.

Figure-17 Semantic technology lifecycle return on investment

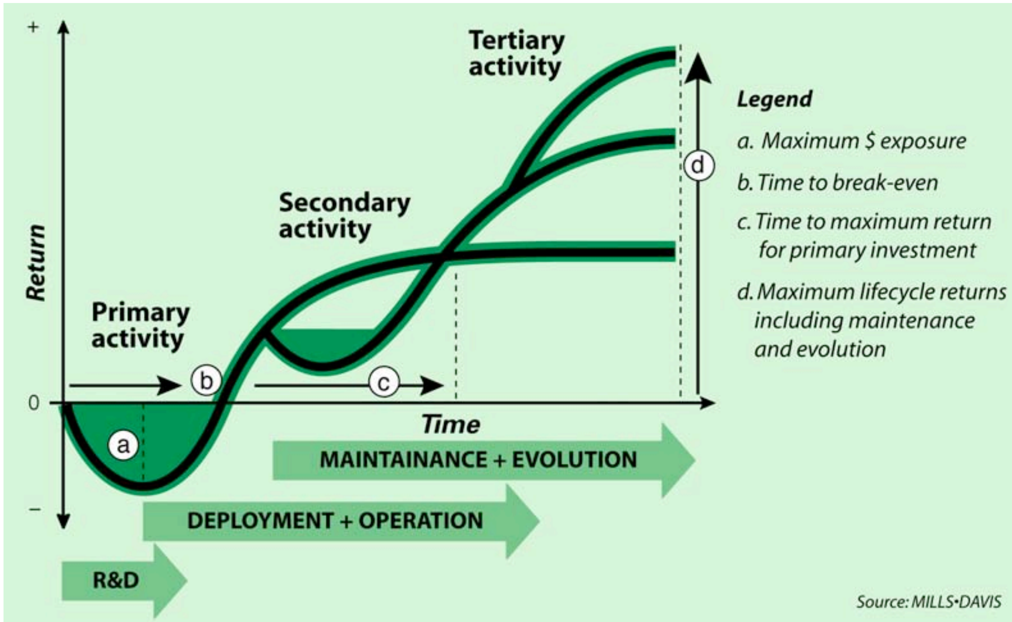

How do semantic technologies maximize lifecycle ROI?

When we map the expenditure of resources versus positive returns across the lifecycle of an enterprise solution, the result is a return-on-investment (ROI) curve (e.g. – revenue and other positive benefits minus capital investment and operating expenses at various points in time, discounted for risk).

Figure-17 depicts hypothetical ROI curves that show cumulative returns across the life of a semantic solution. Time flows from left to right. The amount of return, positive or negative, plots vertically. Stages of the semantic solution lifecycle show as arrows across the bottom of the diagram.

Obviously, the most desirable ROI curve is one that (a) requires zero capital investment, (b) begins to produce positive returns almost immediately, (c) produces substantial positive returns in a reasonable timeframe, and (d) speeds up and drives down the cost of additional related projects that produce even greater returns.

How do semantic technologies maximize value and economic return across each stage of the investment lifecycle?

R&D

The first part of the curve depicts the innovation stage. Every project begins at zero, with nothing ventured and nothing gained or lost. The slope of the graph during innovation is negative. Investments outweigh returns. This is the time of greatest risk and greatest exposure. A key question in evaluating this stage of a project is: how deep into its pockets is the enterprise being asked to go?

A semantic wave approach impacts R&D stage ROI as follows:

- Knowledge needs modeling allows early validation and iterative refinement of requirements, minimizing cost and risk.
- Semantic modeling of UI, data, and system interrelationships minimizes time/cost to prototype.
- Semantic modeling facilitates switching between make, buy, rent, share options at least cost.
- Reduced coding minimizes labor, time, and cost for interoperability, integration, federation.
- Semantic models and composite applications provide unified UI across multiple legacy systems, services and data sources, preserve legacy value, minimize disruption to operations, and reduce development, training, use, and maintenance costs.
- Fast, incremental, non-invasive development cycles accelerate time to value, reduce cost to solution and mitigate development risk.

Deployment and operations

The second part of the curve depicts the operations stage. Solution deployment and initial operations frequently overlap. As operations phase in, the slope of the curve slows its decent, levels and begins to rise. This is called the inflection point of the curve. Returns accumulate going forward. When benefits exceed operating outlays the curve turns positive.

When cumulative returns equal cumulative investments — this is the break-even point. If the time to break-even takes too long, the project may be a bad risk. The curve continues to rise ("in the black") so long as benefits such as revenues exceed operating costs. Net present value analysis is used to compare the relative return on assets employed. Eventually, the benefit stream will slow. Eventually, requirements change, the curve flattens, and the project reaches a point of diminishing returns. Key questions in evaluating this stage of a project are: how long will the company have to wait for positive returns? Also, what is the maximum positive benefit (or upside)?

A semantic wave approach impacts deployment and operations stage ROI as follows:

- Semantic solutions deploy rapidly, incrementally, iteratively, and flexibly, resulting in lower exposure and faster time to value.
- During operation, the semantic model integrated solutions require the less overhead for staffing and support, which helps reduce total cost of ownership.

- Composite applications provide common context and access to underlying information and processes so that users do not have to learn multiple methods to search and navigate across them, which increases their productivity.
- Semantic model driven solutions can be self-documenting and self-explaining, which reduces training and support costs, and helps mitigate risks from knowledge erosion when personnel change roles.
- Semantic models make security and robustness of the deployment much easier (and less expensive) to ensure for mission critical workflows.

Maintenance and evolution

The third part of the curve depicts secondary and tertiary maintenance and enhancement projects that build off of the solution established by the primary project. The measure of performance that is relevant here is the ratio of added value to added cost and risk. A good ROI curve would enable these projects to begin in a timely manner, and be funded by positive returns from the base project. A key question in evaluating this stage of a project is: What is the total upside for related projects that can be funded from the proceeds of this project?

A semantic wave approach impacts maintenance and evolution stage ROI as follows:

- Semantically modeled solutions are easier to scale up and scale out — adding new capabilities, users, locations, security or capacity.
- Semantic models and open standards (knowledge plane) insulate components to minimize impact of changes. This facilitates best-of-breed substitutions, integration of new capabilities, and extension to embrace legacy applications. Faster time to enhance, lower switching costs.
- Semantic models provide leverage to accelerate secondary and tertiary ROIs. Less capital re-investment, less development risks.

Figure-18 Performance of semantic technologies

EFFICIENCY	EFFECTIVENESS	EDGE
Cost savings	*Return on assets*	*Return on investment*
Doing the same job faster, cheaper, or with fewer resources than it was done before	Doing a better job than the one you did before, making other resources more productive and increasing their return on assets and attainment of mission	Changing some aspect of what the business does, resulting in growth, new value capture, mitigation of business risk, or other strategic advantage
EARLY ADOPTER CASE EXAMPLES		
20-80% less labor hours 20-90% less cycle time 30-60% less inventory levels 20-75% less operating cost 25-80% less set-up & development time 20-85% less development cost	50-500% quality gain 2-50X productivity gain 2-10X greater number or complexity of concurrent projects, product releases & units of work handled 2-25X increased return on assets.	2-30X revenue growth 20-80% reduction in total cost of ownership 3-12 month positive return on investment 2-300X positive ROI over 3-years

Source: MILLS•DAVIS

How do semantic technologies improve performance?

The classic motivations for investing in new technologies are basically three:

- Efficiency gain
- Effectiveness gain
- Strategic edge

Efficiency gains mean doing the same job faster, cheaper, or with fewer resources than it was done before. The key measurement is cost savings. Semantic technologies can have a dramatic impact on labor hours, cycle time, inventory levels, operating cost, development time and cost. Early adopter case examples showed 20-90% reductions in these measures.

Effectiveness gains means doing a better job than the one you did before, making other resources more productive, and improving the attainment of mission. The key measurement is return on assets. Semantic technologies can drive dramatic improvements in quality, service levels, and productivity. Combined with process improvements, these can allow existing staff to handle a greater number (or complexity) of current projects, product releases, and units of work. Early adopter case examples showed increases in effectiveness and return on assets from 2-50 times.

Strategic edge means changing some aspect of what the entity does, resulting in growth, new value capture, mitigation of business risk, or other strategic advantage. The key measurement is return on investment. The strategic value of semantic technologies comes from new capabilities that tap new sources of value, resulting in new advantages.

5. Signs of Life in the Marketplace

Figure-19 Semantic wave R&D and early adopter case examples

Where, how, and in what ways do semantic technologies have application?

While early adopter experiences are never uniformly positive, the number reporting positive outcomes, the diversity of applications, and the range of industry verticals represented encourage us. The semantic wave is building, and now appears poised to "cross the chasm" from early adoption to mainstream markets.

Figure-19 depicts areas in business and government where semantic technologies are being applied and delivering value today. It is based on more than 100 early adopter case examples from different economic sectors, including: government, financial services, manufacturing, logistics, transport and communications, energy, health and life sciences, media, and business services. Semantic solutions surveyed delivered 2-10x improvements in measures of performance across the investment lifecycle.

Listed clockwise from the top, these areas include: (1) managing risk, (2) customer-facing services, (3) output management, (4) "smart" products and services, (5) design and manufacture, (6) research, (7) input management, (8) supplier-facing processes, (9) infrastructure and integration, and (10) intelligence, security, and decision-support.

Early adopter case examples and assessments of best practices and lessons learned will be provided in *Semantic Wave 2006: Part-2.*

Figure-20 Technology providers developing semantic solutions

Active Navigation	CheckMi	Empolis	Intelliseek	Noetix	SilkRoad
Adobe	Cisco	Endeca	Intellisophic	Northrop Grumman	Software AG
Aduna	ClearForest	Engenium	Interwoven	nStein	Sony
Agilense	CoeTruman Technologies	Enigmatec	Inxight	NuTech	SRA International
AKT Triple Store	Cogito	EnLeague Systems	ISX Software	Ontologent	SRI International
Amblit Technologies	CognIT	Entopia	ISYS Search Software	Ontology Works	Stanford University
Anteon	Cognos	Entrieva	JARG	Ontopia	Stellent
Apelon	Composite	Epistemics Ltd.	Jayna	Ontoprise	Stratify
APR Smartlogik	Compoze Software	Factiva	Kalido	OpenText	Sun Microsystems
Arbortext	Computer Associates	Fair Isaac	Kanisa Software	Oracle	Sybase
Ask Jeeves	Conformative Systems	FAST	Knowledge Foundations	Profium	Synomos
AskMe	Connecterra	FileNet	Knowledge Media Institute	Radar Networks	SYS Technologies
Aspasia	Connotate	Fujitsu	Kofax	Raytheon	Tacit
Astoria Software	Content Analyst	GeoReference Online	Kowari	Readware	Taxonomywarehouse
AT&T	Contextware	Global360	L&C	RuleBurst	TEMIS
ATG	Contivo	Gnowsis	Lockheed Martin	Reed Elsivier	The Brain
Attensity	Convera	Google	Logic Library	SAIC	Thetus
Autonomy	Copernic	Grand Central	Mark Logic	Sandpiper Software	Thomson
Axontologic	Correlate	Groxis	McDonald Bradley	SAP	TopQuadrant
BBN	Cougaar Software	H5 Technology	Metacarta	SAS	Triple Hop
BEA	Coveo Solutions	Hewlett Packard	MetaIntegration	SchemaLogic	Troux
BioWisdom	Crystal Semantics	Hummingbird	Metallect	Semagix	Ultimus
Black Pearl	Cycorp	Hyperion	Metamatrix	Semandex Networks	Unicorn
Blue Oxide	Dassault Systems	i2 Inc	Metatomix	Semantic Light	Verity
BrandSoft	DAY	IBM	Microsoft	Semantic Research	Versatile Info Sys
Broadvision	Digital Harbor	iLog	Mind Alliance	Semantic Sciences	VerticalNet
Business Objects	Discovery Machine	Image Matters	Miosoft	Semansys	Vignette
C24 Solutions	Dynamic Digital Media	Informatica	Modulant	Semaview	Visual Knowledge
Capraro Technologies	Dream Factory	InforSense	Mondeca	Semtation GmbH	Vitria
Captiva	EasyAsk	Infosys	NCR Teradata	Serena	Vivisimo
Celcorp	Ektron	Innodata (ISOGEN)	NetMap Analytics	SiberLogic	WiredReach
Cerebra	EMC/Documentum	Intellidimension	Neurok	Siderean	XSB

Source: MILLS•DAVIS—10/01/2005

Who is developing semantic solutions?

Project10X research indicates that nearly 200 business entities are engaged in semantic technology R&D for development of products and services to deliver solutions. More than 70 have announced and launched semantic technology based products or services.

Supplier positioning, profiles and technology category assessments will be provided in *Semantic Wave 2006: Part-2.*

6. Semantic Wave Market View

Figure-21 Semantic wave market structure

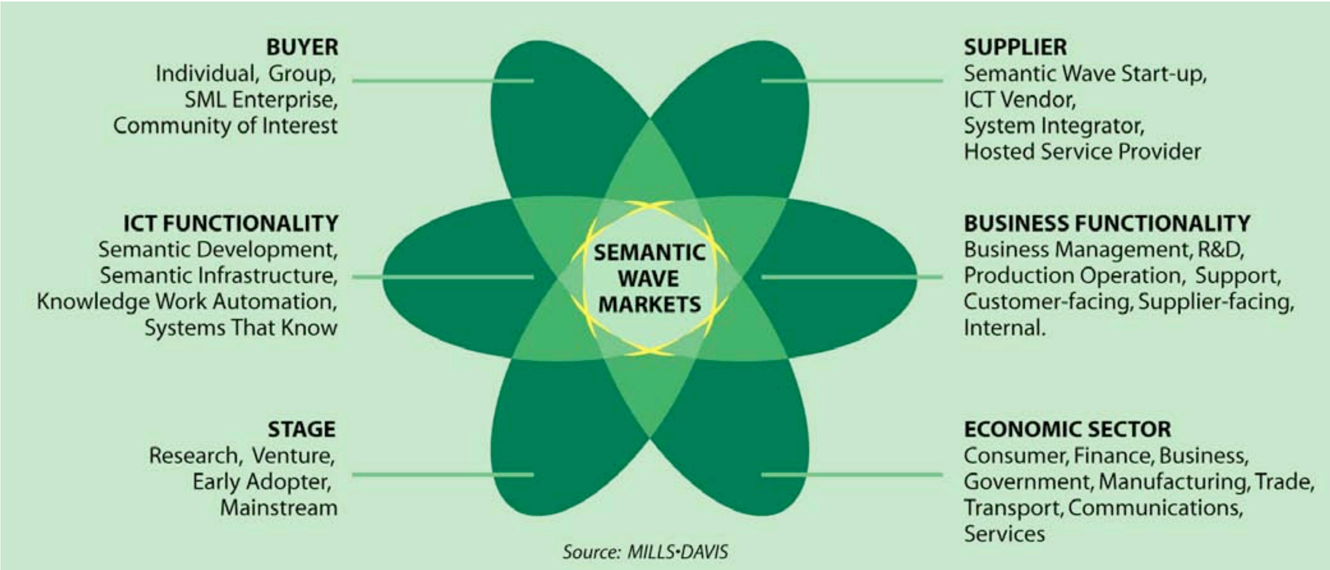

What is the structure of semantic wave markets?

The structure of semantic wave markets can be described in terms of its players (buyers and suppliers), the capabilities provided (ICT and business functionality), stages of development (R&D to mainstream), and sectors of the economy being served.

Figure-21 views semantic wave markets from the perspective of six intersecting and complementary market segments:

- *Buyer* — Who buys semantic technologies? The buyers of semantic wave products and services include: Individuals, Groups, Small, medium and large enterprises, and Communities of interest.
- *Supplier* — Who supplies semantic wave products and services? Providers of semantic wave products and services include: Semantic wave start-ups, ICT vendors, System integrators, Knowledge asset providers, and Hosted service providers.
- *ICT functionality* — What categories of ICT functionality will buyers of semantic wave products and services demand? Buyers of ICT functionality will demand the following categories of semantic capabilities: Semantic development, Semantic infrastructure, Knowledge work automation, and Systems that know.
- *Business functionality* — What categories of business functionality will buyers of semantic wave products and services demand? Buyers of business functionality will demand the following categories of capabilities powered by semantic technologies: Business management, R&D, Production operation, Support, Customer-facing services, Supplier-facing services, Internal services.

- *Stage* — What are the stages of market development for different categories of semantic wave technology and capability? As semantic wave markets develop, technologies and solutions pass through the following stages: Research, Venture, Early adopter, and Mainstream.
- *Economic sector* — What economic sectors are demanding semantic wave capabilities? Buyers in the following economic sectors are demanding semantic wave products and services: Consumer, Finance, Government, Manufacturing, Trade, Transport & Communications, Services, Resources

Market segment profiles and assessments will be provided in *Semantic Wave 2006: Part-2.*

Figure-22 Semantic wave market growth to 2015

MARKET	2006	2010	2015
Semantic Development	\$50M	\$0.4B	\$2.0B
Semantic Infrastructure	\$500M	\$17.0B	\$200.0B
Knowledge Work Automation			
Information-intensive	\$1,100M	\$30.0B	\$250.0B
Knowledge-intensive	\$350M	\$4.5B	\$40.0B
Systems That Know	\$100M	\$0.5B	\$10.0B
TOTAL	\$2,200M	\$52.4B	\$500.0B

Source: MILLS•DAVIS

How will semantic wave markets evolve to 2015?

Markets for semantic technology products and services will grow 10-fold from 2006 to 2010 to more than \$50B worldwide. From 2010 to 2015 the semantic market is expected to grow nearly ten-fold again, fueling trillion-dollar economic expansions worldwide.

Figure-22 depicts overall market growth for semantic ITC functionality segments for semantic development, semantic infrastructure, knowledge work automation, and systems that know.

- *Semantic development* — includes products and services relating to discovery, prototyping, implementation, deploying and maintaining semantic development methodologies, training, tools, platforms, including ontology modeling and life cycle management environments.
- *Semantic infrastructure* — includes products and services relating to semantic enterprise architecture, and semantic technology building blocks for net-centric integration, interoperability, security, such as semantic web services, semantic grid computing, semantic P2P, semantic mobility, semantic technologies for Ipv6, application-oriented networking, service-oriented architecture, semantic business process management, and semantic information interoperability.
- *Information-intensive knowledge work* — includes products and services relating to knowledge work automation through integration and interoperability of systems and information, semantic collaboration, semantic search, semantic-enabled authoring, discovery, and knowledge extraction and categorization, semantic registries and metadata management, composite applications, semantic collaboration, semantic portals.

- *Knowledge-intensive knowledge work* — includes products and services relating to automating knowledge work through knowledge computing such as knowledge asset development and lifecycle management of very large knowledgebases, knowledge commerce and knowledge sharing, knowledge tools for research, analysis, planning, design, diagnosis, simulation and testing, and decision-management.
- *Systems that know* — includes products and services relating to development of cognitive systems, intelligent agents, robots and systems that know what their doing.

Evolution of semantic wave markets

It is helpful to distinguish between where semantic wave markets are today, where they are heading currently, and where they will really go over the next decade.

Where are semantic wave markets today?

Currently, the semantic wave appears as a tiny fraction of ITC markets. Information technologies, stack architecture, and procedural algorithmic programming paradigms dominate. The market and installed base is huge — around $1.2T for hardware, software and services.

Automation of transaction systems, proliferations of PCs, and global access through the worldwide web have been crowning achievements. However, IT approaches are hitting a wall in their ability to handle the explosion of infrastructure, information sources, communities of interest, and knowledge that characterizes todays net-centric global marketplace. Integration and interoperability of systems and information have become a large, costly, intractable problem.

Where are ICT markets headed in the near-to-mid-term?

ITC markets are moving towards the semantic web and related open standards based technologies to help to solve some of these information and system plumbing problems. Rather than new markets and product categories, most semantic web applications are a metamorphosis or extension of existing categories. XML is moving towards RDF and OWL. The National Institute for Standards and Technology (NIST) has stated that all standards should be specified as ontologies. Web services are on their way to becoming semantic web services. Grid computing is becoming semantic grid. Search becomes semantic search. Wikis and Blogs are becoming semantic Wikis and semantic Blogs. And, so on.

There is a driving force for this stage of the market. The economics of semantic interoperability for infrastructure- and information-intensive application categories are compelling. By linking systems and information sources together through a shared semantic model or information layer, developers only have to write one interface rather than a myriad of point-to-point interfaces to all the systems and data sources that the model connects to. Solving information and system plumbing issues will deliver significant value, especially for information-intensive knowledge work. Early adopter research reports 2-10X gains in measures of performance over the lifecycle of the investment.

Where will the markets really go over the next decade?

As it is today, the semantic web improves upon IT, but only allows us to garner low-hanging fruit. Much larger and more valuable opportunities exist. First is knowledge worker automation. Second are systems that know, learn and reason as people do.

While currently small (less than $1B), these markets are potentially very large, have dramatic long-term growth potential, and will power trillion dollar economic expansions worldwide by 2015-2020.

Addressable early markets exist whose buyers have knowledge-intensive needs that are unmet.

Early solution models exist that awaited a universal knowledge technology that can handle massive amounts and kinds of knowledge; can reason natively with uncertainty, logic, and values of unlimited complexity; and can effect robust adaptive, autonomous, and learning behaviors. Requirements for knowledge-intensive applications are emerging in the sciences, engineering, professions, management, and entertainment fields. Our market research indicates that the value drivers include new capabilities tapping new sources of value combined with breakthroughs in both effectiveness and efficiency. This will enable solutions delivering up to 100X gains in performance over contemporary approaches.

Conclusion

Semantic wave is a fundamental shift in paradigm, technology, and economics. The entire ICT stack can be viewed as getting an knowledge backplane. The scale, complexity, functionality, and performance requirements of net-centricity, knowledge worker automation, and autonomous systems make this inevitable. *All aspects of ICT will be affected by semantic technologies!*

W3C semantic web proponents are emphasizing “it's about the data”, that is, the role of RDF & OWL in exposing data of all kinds across the web via metadata and subject ontologies, acknowledging that many issues of reasoning and trust are yet to be resolved, and avoiding any mention of “artificial intelligence”. The potential impact of Internet-wide data sharing via semantic web technologies is massive.

OMG committees, on the other hand, are focusing on semantics for executable (model-driven) architecture. This is valuable for systems and information plumbing, and for accelerating the software development cycle. On the other hand, the approach is too limited to aid reasoning over content.

Most ICT analysts are still analyzing the market in terms of the traditional IT stack and its current product categories. This is OK, but it gives too incremental and fragmented a view of what is happening.

Project10X projections to 2010 and beyond are based on a more comprehensive assessment of the issues of infrastructure and information interoperability, knowledge work automation, and intelligent system behaviors.

Issues of interoperability and integration require syntactical, structural, and semantic solution. Syntax gives us the formats for message exchange. Structure gives us the entities and attributes of records and schemas. Semantics, however, are needed to enable us to understand what the information means in our context of use.

System and information plumbing issues cannot be solved in any sustainable fashion without semantic technologies. We cannot cope with the possibilities of IPv6 without semantics. For SOA to be scalable and viable, we need semantics that enable machines to manage the stack and the service interactions across it. Similarly, web services must become semantic web services if we are to significantly reduce the amount of programmer labor required. The same is true for BPM, EAI, EII, Grid, P2P, RFID, security, and mobility solutions. Context and connectivity demand semantics. In fact, any time we're talking system-of-systems, virtual "whatever", composite applications, collaboration, then we are talking semantics.

Problems of knowledge worker automation cannot be solved adequately without semantics. Interoperability in this context means we need to be able to discover, access, understand, and trust information derived from multiple sources in a specific context of use. Yes, we can search Google and find a starting point. But, in many enterprise contexts, this is simply not enough. If we're involved in anything that is enterprise critical, then our needs quickly move from recovery, to discovery, to intelligence gathering, to question answering, to reasoning over the theory and information contained in content and data. The path we're on is from search to knowing. Executable knowledge integrates theory and information derived from many sources across the net.

Interoperability has a social dimension as well. Viewed from the perspective of Web 2.0, problems of interoperability are essentially issues of the quality of the user experience, and the answer entails semantics in the user interface. People want to discover, access, organize, utilize whatever, with the most satisfaction and minimum effort on their part. If it's a cell phone, they've limited screen real estate (& bandwidth), and even more limited attention spans. People need semantics to organize the context of the experience, as well as provide access to the information and services.

About MILLS•DAVIS

Mills Davis

The author of this report is Mills Davis. He is Project10X's managing director for industry research and strategic programs. Mills consults with technology manufacturers, global 2000 corporations, and government agencies on next-wave semantic technologies and solutions.

Mills serves as lead for the Federal CIO council's Semantic Interoperability Community of Practice (SICoP) research into the business value of semantic technologies. Also, he a founding member of the AIIM interoperable enterprise content management (iECM) working group, and a founding member of the National Center for Ontology Research (NCOR).

A noted researcher and industry analyst, Mills has authored more than 100 reports, whitepapers, articles, and industry studies.

Project10X

Project10X is an industry research, education, and consulting initiative that champions the business value of semantic wave technologies and solutions.

Project10X business consultants and technologists bring expertise in strategy development and business modeling, artificial intelligence, object technology, adaptive systems, ontology engineering, knowledge and content management, publishing and media, semantic web and grid computing, and methodologies for knowledge, software and systems engineering.

We apply semantic platforms, tools, methodologies, and services to help enterprises envision, architect, plan and realize high-performance business practices and semantically integrated processes while sustaining and optimizing their investments in current technologies.

Project10X's goal is to help clients achieve capabilities of strategic value, major cost reduction and cycle time improvement in mission critical workflow, added value for customers, hi-yield capital investment, rapid ROI, and greatly reduced life cycle risk.

